arXiv:1507.05593v1 [cs.NA] 20 Jul 2015

AN EFFICIENT SOLVER FOR SPARSE LINEAR SYSTEMS BASED ON RANK-STRUCTURED CHOLESKY FACTORIZATION

JEFFREY N. CHADWICK AND DAVID S. BINDEL

Abstract. Direct factorization methods for the solution of large, sparse linear systems that arise from PDE discretizations are robust, but typically show poor time and memory scalability for large systems. In this paper, we describe an efficient sparse, rank-structured Cholesky algorithm for solution of the positive definite linear system $\mathbf{Ax} = \mathbf{b}$ when $\mathbf{A}$ comes from a discretized partial-differential equation. Our approach combines the efficient memory access patterns of conventional supernodal Cholesky algorithms with the memory efficiency of rank-structured direct solvers. For several test problems arising from PDE discretizations, our method takes less memory than standard sparse Cholesky solvers and less wall-clock time than standard preconditioned iterations.

Key words. supernodal Cholesky, preconditioners, low-rank structure, randomized algorithms

AMS subject classifications. 65F05, 65F08, 65F50

1. Introduction.

We consider the problem of solving a sparse linear system

$$\mathbf{Ax} = \mathbf{b} \tag{1.1}$$

in which $\mathbf{A}$ is symmetric and positive definite (SPD). In particular, we consider the Cholesky factorization

$$\mathbf{A} = \mathbf{LL}^T, \tag{1.2}$$

where $\mathbf{L}$ is a sparse lower triangular matrix; see [12, 9, 7]. After factoring $\mathbf{A}$, one solves (1.1) by two triangular solves, at cost proportional to the number of nonzeros in $\mathbf{L}$. This approach solves (1.1) exactly up to roundoff effects, and modern supernodal factorization algorithms achieve high flop rates by organizing the factorization around dense matrix kernels. The chief drawback of sparse direct methods is that the factor $\mathbf{L}$ may generally have many more nonzero elements than $\mathbf{A}$. These *fill* elements limit scalability of the method in both time and memory used, particularly for problems coming from the discretization of three-dimensional PDEs, where the number of nonzeros in $\mathbf{L}$ typically scales as $O(N^{3/2})$, where N is the dimension of $\mathbf{A}$.

Compared to direct factorization, iterative methods for (1.1) generally cost less in memory and in time per step than direct methods, but converge slowly without a good preconditioner. Preconditioning involves a complex balance between the progress in each step and the cost of setting up and applying the preconditioner. Even for a single preconditioner type, there are usually many parameters that are optimized on a problem-by-problem basis. For this reason, packages like PETS provide interfaces to allow users to quickly experiment with different preconditioners and parameter settings [1], while commercial finite element codes often forego the potential benefits of iterative methods and simply use out-of-core direct solvers [26].

Fast direct factorization methods and preconditioned iterative solvers each use a different types of structure. A key idea behind sparse direct methods is that one can use the structure of the graph associated with $\mathbf{A}$ to reason about about fill in $\mathbf{L}$. This graph-theoretic approach underlies many modern sparse matrix algorithms, from methods of computing fill-reducing elimination orderings to "supernodal" factorization methods organized around dense matrix operations on columns with similar nonzero structure [7]. In contrast, to solve problems arising from elliptic PDE discretization efficiently, multi-level preconditioners exploit the elliptic regularity of the underlying differential equation. Building on ideas from fast direct

solvers for integral equations [15], recent work in "data-sparse" direct solvers uses both types of structure at once, computing a structured factorization that incorporates (approximate) low-rank blocks [14, 30, 25, 28, 13].

In this paper, we describe an efficient sparse, rank-structured Cholesky algorithm for the solution of (1.1) when $\mathbf{A}$ comes from discretization of a PDE. Our method combines the efficient memory access patterns of conventional supernodal Cholesky algorithms with rank-structured direct solvers. Unlike prior solvers, our method works as a "black box" solver, and does not require information about an underlying PDE mesh. For several test problems arising from PDE discretizations, we show that our method takes less memory than standard sparse Cholesky codes and wall-clock time than standard preconditioners The remainder of the paper is organized as follows. In Section 2, we briefly review the standard supernodal left-looking sparse Cholesky algorithm on which our method is based. In the supernodal factorization, each supernode has an associated diagonal block storing interactions within that supernode, and an off-diagonal block for storing interactions between supernodes. In Section 3, we describe how our algorithm forms and uses low-rank approximations to the off-diagonal blocks of the supernodes, and in Section 4, we describe our approach to hierarchical compression of the diagonal blocks. We discuss some key implementation details in Section 5, and illustrate the behavior of our algorithm on several example problems in 6. Finally, in Section 7, we conclude and give potential directions for future work.

2. Background and Notation. We focus primarily on *supernodal left-looking Cholesky factorization*. This method has yielded implementations which make effective use of modern computing architectures to efficiently solve (1.1) [6].

2.1. Supernodal Left-Looking Cholesky Factorization. Most sparse Cholesky codes have two phases: a fast symbolic analysis phase to compute the nonzero structure of $\mathbf{L}$, and a more expensive numerical factorization phase in which the actual elements of $\mathbf{L}$ are computed. The symbolic analysis phase is organized around an *elimination tree* that encodes the structure of $\mathbf{L}$: in general, $l_{ij} \neq 0$ precisely when there is some k such that $a_{ik} \neq 0$ and j is reachable from k by an elimination tree path that passes only through nodes with indices less than i. Often, the elimination tree has chains of sequentially-numbered nodes corresponding to columns with similar nonzero structure; these can be seen as *supernodes* in a coarsened version of the elimination tree. Supernodal factorization algorithms organize $\mathbf{L}$ around such supernodes, formed by collecting adjacent columns which are predicted to have similar non-zero patterns in $\mathbf{L}$. Supernodal methods achieve high efficiency by storing the nonzero entries for a supernode together as a dense matrix, and by operating on that matrix with optimized kernels from the BLAS.

Suppose $\mathbf{L} \in \mathbb{R}^{N \times N}$ is partitioned into M supernodes. Let $\mathcal{C}_j = (c : c_j \leq c < c_{j+1})$ refer to the column indices in supernode j, and let $\mathcal{C}_j^O = (c : c \geq c_{j+1})$ be the list of columns occurring *after* supernode j. Finally, let $\mathbf{L}_j$ refer to the block column $\mathbf{L}_j = \mathbf{L}(:, \mathcal{C}_j)$. Since $\mathbf{L}$ is lower triangular, it follows that $\mathbf{L}_j(1 : c_j - 1, :) = \mathbf{0}$. We store the matrix $\mathbf{L}_j(\mathcal{C}_j, :)$ explicitly as a dense matrix and refer to this as the supernode's *diagonal block* $\mathbf{L}_j^D$. We also define $\mathcal{R}_j$ to be the list of nonzero rows of $\mathbf{L}_j$ below the diagonal block; that is,

$$\mathcal{R}_j = \left(k \in \mathcal{C}^O : \exists p \in \mathcal{C}_j,\, \ell_{k,p} \neq 0\right) = \left(r_j^1, r_j^2, \ldots\right). \tag{2.1}$$

Since columns in $\mathcal{C}_j$ have similar non-zero patterns, we store $\mathbf{L}_j(\mathcal{R}_j, :)$ as a dense matrix and refer to this as supernode j's (compressed) off-diagonal block $\mathbf{L}_j^O$.

Left-looking algorithms such as the one implemented in [6] form block columns $\mathbf{L}_j$ in order from left to right. We identify the *descendants* $\mathbb{D}_j$ of supernode j as follows:

$$\mathbb{D}_j = \{1 \leq k < j : \mathcal{R}_k \cap \mathcal{C}_j \neq \emptyset\}; \tag{2.2}$$

that is, supernodes from earlier in the factorization whose off-diagonal row set intersects the column set of node j. We also refer to node j as an *ancestor* of node k if $k \in \mathbb{D}_j$. For convenience, we also define index lists relating rows in supernode j to rows in a the off-diagonal block $\mathbf{L}_k^O$ of a descendant supernode k:

$$R_{k\to j}^D = (1 \le p \le |\mathcal{R}_k| : r_k^p \in \mathcal{C}_j)\,, \tag{2.3}$$

$$R_{k\to j}^O = (1 \le p \le |\mathcal{R}_k| : r_k^p \in \mathcal{R}_j)\,. \tag{2.4}$$

Intuitively, (2.3) helps us extract the rows of $\mathbf{L}_k^O$ that influence the contents of $\mathbf{L}_j^D$. Similarly, we use (2.4) to extract rows of $\mathbf{L}_k^O$ needed to form $\mathbf{L}_j^O$. We will also write $\mathcal{R}_{k\to j}^D$ to denote the index list corresponding to $R_{k\to j}^D$ in the uncompressed structure, $\left(r_k^p \in \mathcal{R}_k : p \in R_{k\to j}^D\right)$, and similarly for $\mathcal{R}_{k\to j}^O$.

We also find it convenient to define the function $\mathrm{scatterRows}(\mathbf{B}, R_1, R_2)$, where $R_1 \subseteq R_2$ are ordered index lists, and $\mathbf{B}$ has $|R_1|$ rows. This function returns a matrix with $|R_2|$ rows by placing the contents of rows of $\mathbf{B}$ in the output according to the positions of entries of R_1 in R_2. For example,

$$\mathrm{scatterRows}\left(\begin{pmatrix} 1 & 2 \\ 3 & 4 \end{pmatrix}, \{3,8\}, \{2,3,5,8\}\right) = \begin{pmatrix} 0 & 0 \\ 1 & 2 \\ 0 & 0 \\ 3 & 4 \end{pmatrix}. \tag{2.5}$$

We similarly define the functions $\mathrm{scatterColumns}(\mathbf{B}, C_1, C_2)$, and $\mathrm{scatter}(\mathbf{B}, R_1, R_2, C_1, C_2)$ which composes scatterRows and scatterColumns. Finally, we define gatherRows as a function which reverses the operation of scatterRows; e.g., if $\mathrm{scatter}(\mathbf{B}, R_1, R_2) = \mathcal{C}$ then $\mathrm{gatherRows}(\mathcal{C}, R_2, R_1) = \mathbf{B}$.

Forming the numerical contents of supernode j begins with assembly of a block column of the Schur complement:

$$\mathcal{U}_j^D = \mathbf{A}(\mathcal{C}_j, \mathcal{C}_j) - \sum_{k\in\mathbb{D}_j} \mathrm{scatter}\left(\mathbf{L}_k^O\left(R_{k\to j}^D, :\right)\mathbf{L}_k^O\left(R_{k\to j}^D, :\right)^T,\, \mathcal{R}_{k\to j}^D,\, \mathcal{C}_j,\, \mathcal{R}_{k\to j}^D,\, \mathcal{C}_j\right) \tag{2.6}$$

$$\mathcal{U}_j^O = \mathbf{A}(\mathcal{R}_j, \mathcal{C}_j) - \sum_{k\in\mathbb{D}_j} \mathrm{scatter}\left(\mathbf{L}_k^O\left(R_{k\to j}^O, :\right)\mathbf{L}_k^O\left(R_{k\to j}^D, :\right)^T,\, \mathcal{R}_{k\to j}^O,\, \mathcal{R}_j,\, \mathcal{R}_{k\to j}^D,\, \mathcal{C}_j\right) \tag{2.7}$$

$\mathcal{U}_j^D$ and $\mathcal{U}_j^O$ are dense matrices with the same sizes as $\mathbf{L}_j^D$ and $\mathbf{L}_j^O$. We note that if node k is a descendant of node j, then $\mathcal{R}_k \cap \mathcal{C}_j^O \subseteq \mathcal{R}_j$. The Schur complement in node j is formed by first extracting dense row subsets of $\mathbf{L}_k^O$ for each descendant k, then forming the matrix products from (2.6-2.7) using dense matrix arithmetic and finally scattering the result to $\mathcal{U}_j^D$ and $\mathcal{U}_j^O$. Next, the diagonal and off-diagonal blocks of supernode j are formed as follows:

$$\mathbf{L}_j^D = \mathrm{chol}\left(\mathcal{U}_j^D\right) \tag{2.8}$$

$$\mathbf{L}_j^O = \mathcal{U}_j^O\left(\mathbf{L}_j^D\right)^{-T} \tag{2.9}$$

This procedure is summarized in Algorithm 1.

2.2. Fill-Reducing Ordering. Figure 2.1 provides an example of how a nested dissection ordering might be used on a simple two-dimensional domain, and how this ordering influences fill in the Cholesky factor of the reordered matrix.

Algorithm 1: factorSupernode: Computes the diagonal and off-diagonal blocks $\mathbf{L}_j^D$ and $\mathbf{L}_j^O$ for supernode j. Provided inputs are the matrix to be factored, as well as the partially constructed factor $\mathbf{L}$. every supernode $k \in \mathbb{D}_j$ (node j's descendants) is assumed to have already been factored.

input : $\mathbf{A}, j, \mathbf{L}, \mathbb{D}_j$
output: $\mathbf{L}_j^D, \mathbf{L}_j^O$

begin
  // Initialize the Schur complement blocks
  $\mathcal{U}_j^D \longleftarrow \mathbf{A}(\mathcal{C}_j, \mathcal{C}_j)$
  $\mathcal{U}_j^D \longleftarrow \mathbf{A}(\mathcal{R}_j, \mathcal{C}_j)$
  for each $k \in \mathbb{D}_j$ **do**
    // Build dense update blocks
    diagUpdate $\longleftarrow \mathbf{L}_k^O(R_{k\to j}^D, :) * \mathbf{L}_k^O(R_{k\to j}^D, :)^T$
    offDiagUpdate $\longleftarrow \mathbf{L}_k^O(R_{k\to j}^O, :) * \mathbf{L}_k^O(R_{k\to j}^D, :)^T$
    // Scatter updates to the Schur complement
    $\mathcal{U}_j^D \longleftarrow \mathcal{U}_j^D - \text{scatter}(\text{diagUpdate}, \mathcal{R}_{k\to j}^D, \mathcal{C}_j, \mathcal{R}_{k\to j}^D, \mathcal{C}_j)$
    $\mathcal{U}_j^O \longleftarrow \mathcal{U}_j^O - \text{scatter}(\text{offDiagUpdate}, \mathcal{R}_{k\to j}^O, \mathcal{R}_j, \mathcal{R}_{k\to j}^D, \mathcal{C}_j)$
  // Factor node j's diagonal block
  $\mathbf{L}_j^D \longleftarrow \text{cholesky}(\mathcal{U}_j^D)$
  // Dense triangular solve
  $\mathbf{L}_j^O \longleftarrow \mathcal{U}_j^O(\mathbf{L}_j^D)^{-T}$
  return $\mathbf{L}_j^D, \mathbf{L}_j^O$

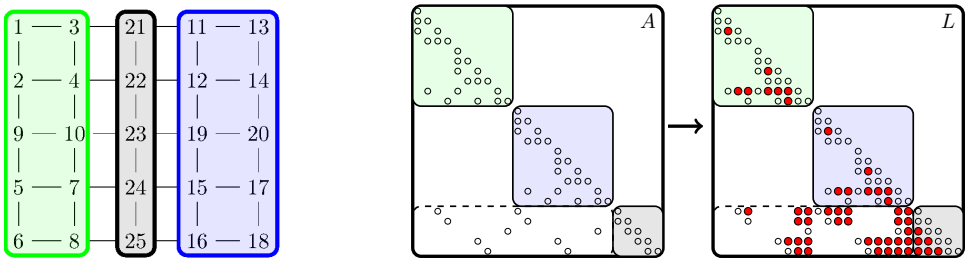

FIG. 2.1. *Fill during nested-dissection Cholesky factorization for a 2D mesh example (left). When nodes in a mesh are ordered so that a vertex separator appears last after the subdomains it separates (right), fill in the Cholesky factor (indicated by red circles) is restricted to the diagonal blocks corresponding to interactions within each subdomain and within the separator and to off-diagonal blocks associated with subdomain-separator interactions.*

3. Off-Diagonal Block Compression. In §2.1 we introduced the notion of supernodes with dense diagonal and off-diagonal blocks $\mathbf{L}_j^D$ and $\mathbf{L}_j^O$. In this section, we discuss the process of approximating off-diagonal blocks $\mathbf{L}_j^O$ with low-rank matrices. The matrix $\mathbf{L}_j^O$ stores interactions between supernode j, and other supernodes in occurring later in the factorization.

3.1. Block Selection and Ordering. We use nested dissection [11] to construct a fill-reducing ordering as originally discussed in §2.2. We choose nested dissection because the geometric structure introduced by this method yields a factor matrix $\mathbf{L}$ in which many dense

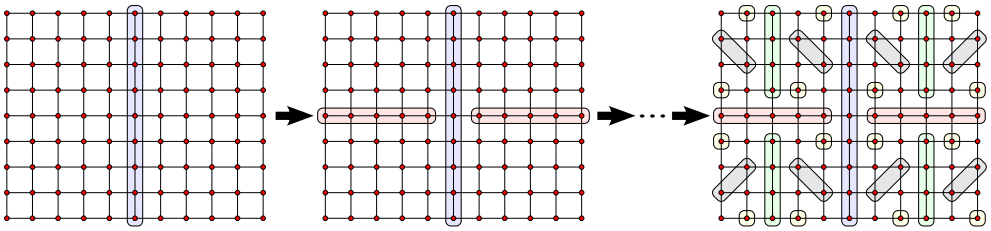

FIG. 3.1. *Nested dissection is applied to a regular, two-dimensional grid. At each level, the domain is recursively subdivided by the introduction of separators. Each image in the sequence depicts a level in the recursive dissection of the domain. Separators at the same level in the nested dissection hierarchy are rendered with the same color. In the rightmost figure we see that five levels of nested dissection fully decompose this domain in to subdomains of unit size.*

submatrices are amenable to low-rank approximation. We will discuss this property in more detail in §3.3. Henceforth, we will assume that the matrix $\mathbf{A}$ has already been symmetrically permuted using nested dissection. We expect that the interactions between large separators in the factorization will exhibit rapidly decaying rank structure (see §3.3). Therefore, we represent each "large" separator from the nested dissection hierarchy with a supernode. The off-diagonal blocks $\mathbf{L}_j^O$ associated with these supernodes describe interactions between large separators in the factor (see Figure 3.2).

Standard sparse Cholesky solvers such as CHOLMOD [6] may optionally use nested dissection for reordering. However, the process of constructing supernodes used by these solvers does not guarantee a one-to-one relationship between supernodes and separators from the nested dissection ordering. In our algorithm, we choose a tolerance $\tau_O \in \mathbb{Z}^+$ and introduce a supernode for every separator with at least τ_O variables. The remaining indices in our reordered matrix $\mathbf{A}$ are gathered in to supernodes using methods identical to [6].

3.2. Block Compression Scheme. Consider the state of the factorization immediately prior to forming the factor contents for supernode j:

$$\begin{pmatrix} \mathbf{A}_{pre}^D & & sym \\ \mathbf{A}_{pre}^O & \mathbf{A}_j^D & sym \\ & \mathbf{A}_j^O & \mathbf{A}_{post} \end{pmatrix} = \begin{pmatrix} \mathbf{L}_{pre}^D & & \\ \mathbf{L}_{pre}^O & \mathbf{I} & \\ & \mathbf{0} & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{L}_{pre}^D & & \\ \mathbf{L}_{pre}^O & \mathbf{U}_j^D & \\ & \mathbf{U}_j^O & \mathbf{U}_{post} \end{pmatrix}^T \tag{3.1}$$

Here *pre* and *post* refer to the sets of columns occurring before and after supernode j, respectively. Note that the Schur complement $\mathbf{U}_{post}$ is never formed explicitly since we only form Schur complements one supernode at a time. We also note that given the definition of $\mathbf{U}_j$ and $\mathbf{L}_j$, it is necessary to apply the $\mathrm{scatter}$ operator to these matrices to make (3.1) valid, but this has been omitted here for brevity. Following factorization of node j, we have:

$$\begin{pmatrix} \mathbf{A}_{pre}^D & sym & sym \\ \mathbf{A}_{pre}^O & \mathbf{A}_j^D & sym \\ & \mathbf{A}_j^O & \mathbf{A}_{post} \end{pmatrix} = \begin{pmatrix} \mathbf{L}_{pre}^D & & \\ \mathbf{L}_{pre}^O & \mathbf{L}_j^D & \\ & \mathbf{L}_j^O & \mathbf{I} \end{pmatrix} \begin{pmatrix} \mathbf{L}_{pre}^D & & \\ \mathbf{L}_{pre}^O & \mathbf{L}_j^D & \\ & \mathbf{L}_j^O & \widetilde{\mathbf{U}}_{post} \end{pmatrix}^T \tag{3.2}$$

where

$$\widetilde{\mathbf{U}}_{post} = \mathbf{U}_{post} - \mathbf{L}_j^O \left(\mathbf{L}_j^O\right)^T. \tag{3.3}$$

Assuming $\mathbf{A}$ is positive definite, $\widetilde{\mathbf{U}}_{post}$ must also be positive definite.

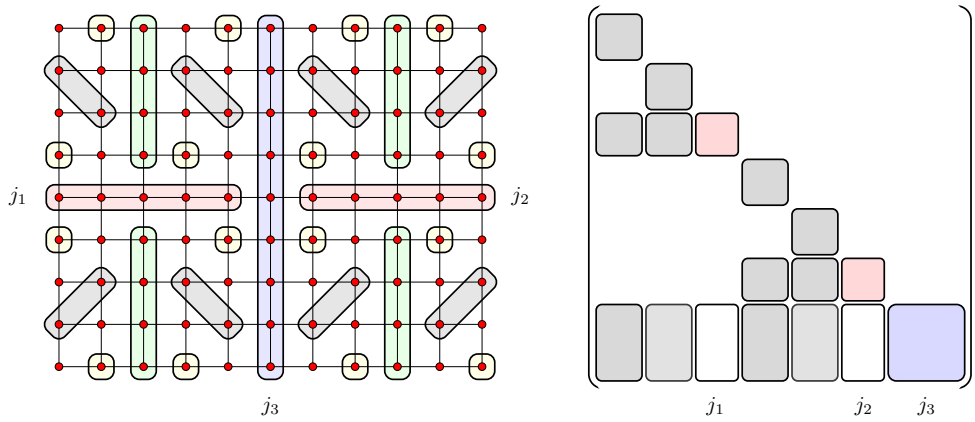

FIG. 3.2. *Consider the two-dimensional grid on the left, decomposed via nested dissection (see Figure 3.1). If we set $\tau_O = 5$, then the three largest separators will be identified as supernodes (labelled j_1, j_2 and j_3). The resulting matrix structure is shown on the right. The off-diagonal blocks $\mathbf{L}^O_{j_1}$ and $\mathbf{L}^O_{j_2}$ are shown in white and describe interactions between separator j_3 and separators j_1 and j_2, respectively. In the domain picture on the left we see that, due to the way these separators intersect geometrically, these interactions tend to mostly occur over large distances. This justifies the use of low-rank matrices to approximate these interactions. The parts of the factor for which we do not apply any compression are shown in gray. These blocks can be evaluated either using the standard supernodal factorization algorithm, or using the interior blocks approach discussed in §5.1.*

If we approximate the off-diagonal part of supernode j with a low-rank matrix – $\mathbf{L}^O_j \approx \mathbf{V}\mathbf{U}^T$ – then the Schur complement in (3.3) is approximated by

$$\overline{\mathcal{U}}_{post} = \mathcal{U}_{post} - \mathbf{V}\mathbf{U}^T\mathbf{U}\mathbf{V}^T. \tag{3.4}$$

We choose $\mathbf{V}$ and $\mathbf{U}$ using a method similar to [23] so that this modified Schur complement (3.4) is guaranteed to remain positive definite. Namely, we choose $\mathbf{U}$ to have orthonormal columns and $\mathbf{V}$ to be the projection of $\mathbf{L}^O_j$ on to this basis; $\mathbf{V} = \mathbf{L}^O_j\mathbf{U}$. We can write $\mathbf{L}^O_j = [\mathbf{V}\ \ \overline{\mathbf{V}}][\mathbf{U}\ \ \overline{\mathbf{U}}]^T$ where $\overline{\mathbf{U}}$ is a (non-unique) matrix with orthonormal columns, $\mathbf{U}^T\,\overline{\mathbf{U}} = \mathbf{0}$, and $\overline{\mathbf{V}} = \mathbf{L}^O_j\overline{\mathbf{U}}$. Given these properties, we can rewrite (3.3) as

$$\widetilde{\mathcal{U}}_{post} = \mathcal{U}_{post} - \mathbf{V}\mathbf{V}^T - \overline{\mathbf{V}}\,\overline{\mathbf{V}}^T \tag{3.5}$$

and (3.4) as

$$\overline{\mathcal{U}}_{post} = \mathcal{U}_{post} - \mathbf{V}\mathbf{V}^T \tag{3.6}$$

$$= \widetilde{\mathcal{U}}_{post} + \overline{\mathbf{V}}\,\overline{\mathbf{V}}^T. \tag{3.7}$$

Since $\widetilde{\mathcal{U}}_{post}$ is positive definite and $\overline{\mathbf{V}}\,\overline{\mathbf{V}}^T$ is positive semi-definite, it follows from (3.7) that $\overline{\mathcal{U}}_{post}$ remains positive definite under this approximation.

3.3. Low-Rank Structure. Recall from §3.1 that we use nested dissection to reduce fill and represent large separators as supernodes in the factorization. Nested dissection orderings are built entirely based on $\mathbf{A}$'s graph structure. However, in problems defined on physical domains (say, discretizations of partial differential equations on two- or three-dimensional domains) separators also have a convenient geometric interpretation. In these problems, separators are geometric regions which bisect subdomains of the original problem domain (see Figures 3.1 and 3.2). For example, in many three-dimensional problems the nested dissection

separators are surfaces which cut the domain in to disjoint pieces. Given a partial factorization of a matrix $\mathbf{A}$, the remaining Schur complement $\mathcal{U}$ behaves like a discretization of a boundary integral equation [5]. For many problems, these discretizations will yield smooth coefficients for matrix indices which are geometrically distant from each other in the original problem domain. The structure of separators produced by nested dissection tends to ensure that the interactions between pairs of large separators occur mostly over large distances, with only a handful of "near-field" interactions (see Figure 3.2). If supernode j corresponds to a large separator and $\mathcal{U}_j^O$ is this node's off-diagonal block in the Schur complement, then we expect that $\mathcal{U}_j^O$ should have rapidly decaying rank structure due to the property discussed above. As such, $\mathcal{U}_j^O$ (and, likewise, $\mathbf{L}_j^O$) admits a low-rank approximation. Therefore, we compress the off-diagonal blocks in supernodes/separators which are sufficiently large (larger than τ_O). We use the following notation for this low-rank approximation:

$$\mathbf{L}_j^O \approx \mathbf{V}_j \mathbf{U}_j^T \quad \text{where} \quad \mathbf{V}_j \in \mathbb{R}^{|\mathcal{R}_j| \times q}, \mathbf{U}_j \in \mathbb{R}^{|\mathcal{C}_j| \times q} \quad \text{and} \quad q \ll |\mathcal{R}_j|, |\mathcal{C}_j| \tag{3.8}$$

Algorithm 2: offDiagonalMultiply: Computes the product $\mathbf{B} = \mathbf{L}_j^O \mathbf{G}$ given some input matrix $\mathbf{G} \in \mathbb{R}^{|\mathcal{C}_j| \times r}, r > 0$. The function $\mathrm{diagonalSolve}(j, \mathbf{X}, \mathsf{transpose})$ applies the inverse of $\mathbf{L}_j^D$ to the input matrix. If transpose is set to true, then diagonalSolve forms the product $(\mathbf{L}_j^D)^{-T}\mathbf{X}$ instead. See §4 and Algorithm 4 for a detailed description of diagonalSolve.

input : $\mathbf{A}$, j, $\mathbf{L}$, $\mathbb{D}_j$, $\mathbf{G}$
output: $\mathbf{B} = \mathbf{L}_j^O \mathbf{G}$

begin
  // Apply node j's diagonal inverse to the input
  $\mathbf{G} \longleftarrow \mathrm{diagonalSolve}(j, \mathbf{G}, \mathsf{transpose} = \mathsf{true})$
  // Multiply by the desired block from $\mathbf{A}$
  $\mathbf{W} \longleftarrow \mathbf{A}(\mathcal{R}_j, \mathcal{C}_j)\mathbf{G}$
  for each $k \in \mathbb{D}_j$ **do**
    // Extract the required sub-matrix from $\mathbf{G}$
    $\mathbf{G}_{sub} \longleftarrow \mathrm{gatherRows}(\mathbf{G}, \mathcal{C}_j, \mathcal{R}_{k \to j}^D)$
    // Form the needed product with two multiplications
    $\mathbf{T} \longleftarrow \left[\mathbf{L}_k^O(R_{k \to j}^D, :)\right]^T \mathbf{G}_{sub}$
    $\mathbf{T} \longleftarrow \mathbf{L}_k^O(R_{k \to j}^O, :)\mathbf{T}$
    // Scatter result to the output matrix
    $\mathbf{W} \longleftarrow \mathbf{W} - \mathrm{scatterRows}(\mathbf{T}, \mathcal{R}_{k \to j}^O, \mathcal{R}_j)$
  return $\mathbf{W}$

3.4. Compression Algorithm. Next we discuss our approach for forming the low-rank approximation $\mathbf{L}_j^O \approx \mathbf{V}_j \mathbf{U}_j^T$. Our compression strategy must satisfy two requirements:

1. The off-diagonal block $\mathbf{L}_j^O$ may be expensive to construct and store. Therefore, we wish to build $\mathbf{V}_j$, $\mathbf{U}_j$ without explicitly constructing $\mathbf{L}_j^O$.
2. The original supernode factorization procedure presented in algorithm 1 makes effective use of dense matrix arithmetic, allowing for very efficient implementations [6]. Our compression algorithm should preserve this property.

To satisfy these requirements, we use randomized low-rank approximation algorithms [24, 16]. Similar randomized methods have been used previously in rank-structured sparse solvers [10, 29]. The key insight behind these randomized algorithms is that a "good" rank-q approximation to a matrix $\mathbf{B}$ can be found by considering products of the form $\mathbf{BG}$, where $\mathbf{G}$ is a randomly generated matrix with $q+p$ columns and $p > 0$ is a small oversampling parameter (typically $p \approx 5 - 10$ is suitable). In this case, we consider a rank-q approximation to be good if it is close (in the 2-norm) to the best rank-q approximation provided by $\mathbf{B}$'s singular value decomposition (SVD). If the singular values of $\mathbf{B}$ decay slowly, then obtaining such an approximation may require us to instead form products of the form $\mathbf{C} = (\mathbf{BB}^T)^s\mathbf{BG}$, where $s \geq 0$ is a small number of power iterations. The theory behind these methods states that $\mathbf{C}$ provides a column basis for a low-rank approximation of $\mathbf{B}$ which is close to optimal. Moreover, constructing $\mathbf{C}$ only requires a small number of matrix multiplications involving $\mathbf{B}$ and $\mathbf{B}^T$. Therefore, requirement 1 above is satisfied. We can also perform these multiplications in a way that leverages dense matrix arithmetic similar to Algorithm 1, satisfying requirement 2.

Algorithm 2 efficiently forums products $\mathbf{L}_j^O\mathbf{G}$ for an arbitrary dense matrix $\mathbf{G}$. As discussed earlier, we also require products of the form $\left(\mathbf{L}_j^O\right)^T\mathbf{G}$. These products are formed by the function $\mathrm{offDiagonalMultiplyTranspose}$. This function has similar structure fo Algorithm 2. Finally, Algorithm 3 uses the $\mathrm{offDiagonalMultiply}$ and $\mathrm{offDiagonalMultiplyTranspose}$ functions to form a low-rank approximation $\mathbf{L}_j^O \approx \mathbf{V}_j\mathbf{U}_j^T$ for node j's off-diagonal block, where $\mathbf{U}_j$ is chosen to have orthonormal columns (as discussed in §3.2). We note that Algorithm 2 assumes the matrix $\mathbf{L}_k^O$ is stored explicitly for all descendants $k \in \mathbb{D}_j$. In practice, some of these blocks may also have been assigned low-rank representations $\mathbf{L}_k^O \approx \mathbf{V}_k\mathbf{U}_k^T$. If this is the case, then we replace lines 10-11 in Algorithm 2 with

$$\begin{aligned}
\mathbf{U}_{prod} &\longleftarrow \mathbf{U}_k^T\mathbf{U}_k \\
\mathbf{T} &\longleftarrow \left[\mathbf{V}_k(R_{k\to j}^D,:)\right]^T\mathbf{G}_{sub} \\
\mathbf{T} &\longleftarrow \mathbf{U}_{prod}\mathbf{T} \\
\mathbf{T} &\longleftarrow \mathbf{V}_k(R_{k\to j}^O,:)\mathbf{T}.
\end{aligned} \tag{3.9}$$

4. Diagonal Block Compression. Section §3 discussed the process of constructing a sparse Cholesky factorization in which off-diagonal interactions between large separators are approximated with low-rank matrices. In large, three-dimensional problems, larger separators may include thousands to tens of thousands of variables. For these problems, the compression scheme from §3 can provide a significant reduction in both memory usage over standard factorizations, while still providing a factor that serves as an excellent preconditioner. However, if supernode j is large, then building the dense diagonal matrix $\mathbf{L}_j^D$ may also require significant computation and storage. In this section, we discuss an approach to compressing diagonal blocks $\mathbf{L}_j^D$.

4.1. Low-Rank Structure. In §3 we saw that low-rank behavior in off-diagonal blocks $\mathbf{L}_j^O$ is exposed by the geometric structure of nested dissection. We can reorder variables within a separator to expose similar low-rank structure within diagonal blocks $\mathbf{L}_j^D$. Since all columns in a given supernode are treated as having the same fill pattern, we are can perform this reordering without affecting accuracy, memory usage, or computation time. Consider the top-level separator shown in figure 3.1. Suppose that the indices of vertices in this separator are ordered sequentially from top to bottom and that the 9×9 diagonal block for this separator is written as a 2×2 block matrix (assume, without loss of generality, that the first block

Algorithm 3: approximateOffDiagonal: Builds a low-rank approximation $\mathbf{L}_j^O \approx \mathbf{V}_j\mathbf{U}_j^T$ for node j's off-diagonal block. This algorithm assumes that node j's diagonal block has already been factored. The function $\mathrm{randomMatrix}(m,n)$ generates an $m \times n$ matrix whose entries are drawn from a Gaussian distribution with mean 0 and unit variance. The function $\mathrm{makeOrthonormal}$ returns an orthonormal basis for the column space of its input matrix. This can be accomplished by means of – for example – a QR factorization.

input : $\mathbf{A}$, j, $\mathbf{L}$, $\mathbb{D}_j$, off-diagonal rank s_j, number of power iterations q
output: $\mathbf{V}_j, \mathbf{U}_j$ such that $\mathbf{L}_j^O \approx \mathbf{V}_j\mathbf{U}_j^T$

begin
  // Initialize a random matrix
  $\mathbf{G} \longleftarrow \mathrm{randomMatrix}(|\mathcal{R}_j|, s_j)$
  // Implicitly form the product $\mathbf{L}_j^O\mathbf{G}$
  $\mathbf{G} \longleftarrow \mathrm{offDiagonalMultiplyTranspose}(\mathbf{A}, j, \mathbf{L}, \mathbb{D}_j, \mathbf{G})$
  // Run additional power iterations
  for $i = 1$ *to* q **do**
    $\mathbf{G} \longleftarrow \mathrm{offDiagonalMultiply}(\mathbf{A}, j, \mathbf{L}, \mathbb{D}_j, \mathbf{G})$
    $\mathbf{G} \longleftarrow \mathrm{offDiagonalMultiplyTranspose}(\mathbf{A}, j, \mathbf{L}, \mathbb{D}_j, \mathbf{G})$
  // Extract an orthonormal row basis for $\mathbf{L}_j^O$
  $\mathbf{U}_j \longleftarrow \mathrm{makeOrthonormal}(\mathbf{G})$
  // Compute $\mathbf{V}_j$ by projecting on to $\mathbf{U}_j$
  $\mathbf{V}_j \longleftarrow \mathrm{offDiagonalMultiply}(\mathbf{A}, j, \mathbf{L}, \mathbb{D}_j, \mathbf{U}_j)$
  return $\mathbf{V}_j, \mathbf{U}_j$

row/column has four entries, and that the second has five):

$$\mathbf{L}_j^D = \begin{pmatrix} \mathbf{L}_{11} & \mathbf{0} \\ \mathbf{L}_{21} & \mathbf{L}_{22} \end{pmatrix}$$

As a result of the ordering discussed above, $\mathbf{L}_{21}$ stores interactions between vertices in the top half of the separator with vertices in the bottom half. As we discussed in §3.3, the spatial separation between these groups of variables suggests that $\mathbf{L}_{21}$ can be approximated with a low-rank matrix $\mathbf{L}_{21} \approx \mathbf{V}\mathbf{U}^T$. We can apply this argument recursively to $\mathbf{L}_{11}$ and $\mathbf{L}_{22}$ to achieve further compression.

The example above assumes that the rows/columns of $\mathbf{L}_j^D$ are ordered such that off-diagonal blocks of $\mathbf{L}_j^D$ exhibit low-rank structure. Finding such an ordering is straightforward when $\mathbf{A}$ comes from a PDE discretization on a regular mesh like the one pictured in Figure 3.1. However, obtaining a suitable ordering for general, three-dimensional problems on irregular domains is nontrivial.

Recall from §3.3 that separators in the nested dissection hierarchy are geometric regions partition the original problem domain. In three dimensions, we intuitively expect these separators to look like two-dimensional surfaces inside of the original problem domain. Our solver exposes low-rank structure in $\mathbf{L}_j^D$ by reordering indices within separators so that off-diagonal blocks in $\mathbf{L}_j^D$ describe interactions between spatially separated pieces of the separator region. We begin by assigning a three-dimensional position $\mathbf{x}_i : i \in \mathcal{C}_j$ to each index associated with supernode/separator j. There are many techniques for spatially partitioning the positions $\mathbf{x}_i$.

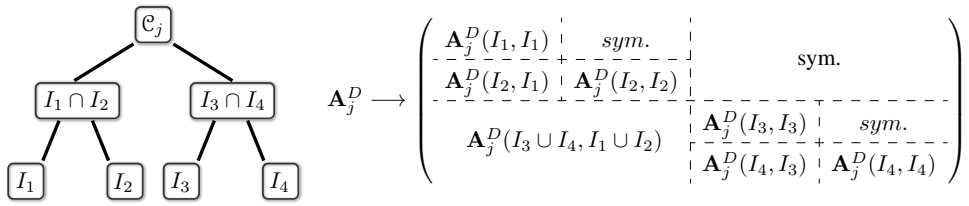

FIG. 4.1. *We recursively partition supernode j's variable indices $\mathcal{C}_j$ resulting in the tree structure shown on the left (assuming two levels of partitioning in this case). We permute the indices in $\mathcal{C}_j$ so that, following this permutation, the index blocks associated with leaves in this tree structure appear sequentially along the diagonal of $\mathbf{A}_j^D$. The resulting permutation to $\mathbf{A}_j^D$ for this two-level example is shown on the right.*

Currently, we use a simple axis-based splitting scheme. We partition the positions $\mathbf{x}_i$ in to two subsets by sorting them along the longest bounding box axis of the set $\{\mathbf{x}_i : i \in \mathcal{C}_j\}$ and splitting this sorted list in to two equal-sized pieces. This process is applied recursively until the separator has been partitioned in to subdomains with at most τ_D variables. The paramter τ_D is chosen in advance as the size of the largest diagonal block that we wish to represent explicitly in the factor matrix. We use this partitioning to reorder the indices within a supernode in a way that exposes low-rank structure. See Figure 4.1 for an explanation of how this permutation is built. Using the two-level partitioning example shown this figure, we label blocks of the diagonal factor block $\mathbf{L}_j^D$ as follows:

$$\mathbf{L}_j^D = \begin{pmatrix} \mathbf{L}_{j,1}^D & \mathbf{0} & \mathbf{0} \\ \mathbf{L}_{j,2}^D & \mathbf{L}_{j,3}^D & \\ \mathbf{L}_{j,4}^D & & \begin{matrix} \mathbf{L}_{j,5}^D & \mathbf{0} \\ \mathbf{L}_{j,6}^D & \mathbf{L}_{j,7}^D \end{matrix} \end{pmatrix}. \tag{4.1}$$

Blocks are numbered in the order in which they must be formed during factorization (intuitively, top to bottom and left to right). In this 4×4 example, $\mathbf{L}_j^D$ is be approximated as follows:

$$\mathbf{L}_j^D \approx \begin{pmatrix} \mathbf{L}_{j,1}^D & \mathbf{0} & \mathbf{0} \\ \mathbf{V}_{j,2}^D (\mathbf{U}_{j,2}^D)^T & \mathbf{L}_{j,3}^D & \\ \mathbf{V}_{j,4}^D (\mathbf{U}_{j,4}^D)^T & & \begin{matrix} \mathbf{L}_{j,5}^D & \mathbf{0} \\ \mathbf{V}_{j,6}^D (\mathbf{U}_{j,6}^D)^T & \mathbf{L}_{j,7}^D \end{matrix} \end{pmatrix}. \tag{4.2}$$

For block s in this matrix, let $R_{j,s}^D$ and $C_{j,s}^D$ be the set of rows and columns over which block s is defined, relative to $\mathbf{L}_j^D$. That is, $\mathbf{L}_j^D(R_{j,s}^D, C_{j,s}^D) = \mathbf{L}_{j,s}^D$. Similarly, let $\mathcal{R}_{j,s}^D$ and $\mathcal{C}_{j,s}^D$ refer to the same row and column sets, but relative to the entire factor $\mathbf{L}$, so that $\mathbf{L}(\mathcal{R}_{j,s}^D, \mathcal{C}_{j,s}^D) = \mathbf{L}_{j,s}^D$.

In summary, we permute the original matrix $\mathbf{A}$ in two main steps. The first is fill-reducing ordering using nested dissection. As we noted in §3, this step exposes low-rank structure in certain off-diagonal submatrices of $\mathbf{L}$, allowing for compression. The second stage of this permutation consists of reordering indices *within* certain supernodes formed in the first stage – namely, those associated with large separators. This does not alter the sparsity of $\mathbf{L}$, but does allow us to compress certain off-diagonal submatrices of these large diagonal blocks.

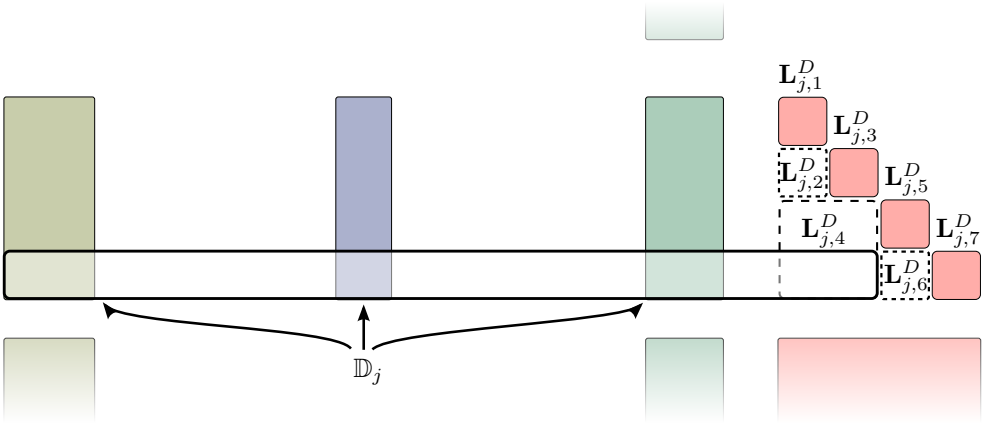

FIG. 4.2. *Block row of a factor in which node j (on the right) has a hierarchically compressed diagonal matrix $\mathbf{L}^D_j$ with structure given by* (4.1)*. Compressed matrices are indicated with a dashed border. Consider the compressed block $\mathbf{L}^D_{j,6}$. When forming products with this matrix (for the purpose of compression), we require contributions from all columns of $\mathbf{L}$ with non-zeros in the region highlighted with a thick black border. Observe that this includes contributions from descendents $\mathbb{D}_j$, as well as other compressed blocks in $\mathbf{L}^D_j$. In this case, $\mathbf{L}^D_{j,6}$ depends on three of node j's descendents, as well as $\mathbf{L}^D_{j,4}$.*

4.2. Compression Algorithm. Next, we turn to compression of off-diagonal blocks within a diagonal matrix $\mathbf{L}^D_j$. As in the compression methods discussed in §3, we use randomized methods to construct low-rank matrix approximations. First, we define the $\mathrm{diagonalSolve}$ function, originally introduced in Algorithm 2. We consider a slight variation on this function, in which the index s of a block from $\mathbf{L}^D_j$ is provided as an argument. Invoking this function with argument s solves a system of equations using the smallest diagonal sub-block of $\mathbf{L}^D_j$ containing $\mathbf{L}^D_{j,s}$. Using (4.2) as an example, calling $\mathrm{diagonalSolve}$ with $s = 2$ would solve a system using the inverse of

$$\begin{pmatrix} \mathbf{L}^D_{j,1} & \mathbf{0} \\ \mathbf{V}^D_{j,2}\left(\mathbf{U}^D_{j,2}\right)^T & \mathbf{L}^D_{j,3} \end{pmatrix}. \tag{4.3}$$

For brevity, calling $\mathrm{diagonalSolve}$ with no "s" argument (as in Algorithm 2) solves systems using the entire matrix $\mathbf{L}^D_j$. This process is summarized in Algorithm 4. Throughout the algorithms discussed in this section, we treat the indices s of blocks in $\mathbf{L}^D_j$ as the labels of nodes in an in-order traversal of a complete binary tree. When we refer to a *child* or *parent* of s, we mean the in-order index of the node which is the child or parent of the node with in-order index s in this tree. For example, if $\mathbf{L}^D_j$ has 7 blocks (as in (4.1)), then $s = 4$ is the root of this tree and has left and right children with indices 3 and 6, respectively.

We will now use the $\mathrm{diagonalSolve}$ function to build an algorithm for compressing an off-diagonal block $\mathbf{L}^D_{j,s}$. As before, we multiply $\mathbf{L}^D_{j,s}$ with random matrices without explicitly constructing $\mathbf{L}^D_{j,s}$ As in §3.4, these operations depend on the contents of node j's descendants $\mathbb{D}_j$. In addition, we may need to consider contributions from previously compressed blocks within $\mathbf{L}^D_j$. See Figure 4.2 for a visual representation of this dependence.

The $\mathrm{diagonalMultiply}$ algorithm (Algorithm 5) provides the details of this procedure.

Algorithm 4: diagonalSolve: Solves a system of equations using the submatrix of $\mathbf{L}_j^D$ rooted at block s in $\mathbf{L}_j^D$'s hierarchical structure.

```
   input : j, G, transpose, s
 1 begin
 2     // Leaf nodes correspond to dense diagonal blocks
 3     if Node s is a leaf then
 4         if transpose then
 5             return (L^D_{j,s})^{-T} G
 6         else
 7             return (L^D_{j,s})^{-1} G
 8     else
 9         // Partition G into two parts
10         G_1 = G(1 : |C^D_{j,s}|, :)
11         G_2 = G(|C^D_{j,s}| + 1 : end, :)
12         // Child blocks of s
13         s_1 = left in-order child of s
14         s_2 = right in-order child of s
15         if transpose then
16             // Backward substitution
17             G_1 ⟵ diagonalSolve(j, G_1, transpose, s_1)
18             G_2 ⟵ G_2 − V^D_{j,s} ((U^D_{j,s})^T G_1)
19             G_2 ⟵ diagonalSolve(j, G_2, transpose, s_2)
20         else
21             // Forward substitution
22             G_2 ⟵ diagonalSolve(j, G_2, transpose, s_2)
23             G_1 ⟵ G_1 − U^D_{j,s} ((V^D_{j,s})^T G_2)
24             G_1 ⟵ diagonalSolve(j, G_1, transpose, s_1)
25         return ( G_1 )
                  ( G_2 )
```

Lines 7-15 in this algorithm resemble the descendant multiplication from Algorithm 2. Lines 17-35 compute contributions from other blocks inside of $\mathbf{L}_j^D$. As before, we also require the algorithm $\mathrm{diagonalMultiplyTranspose}$. This algorithm has a similar structure to Algorithm 5 With these two functions, we define a function $\mathrm{approximateDiagonalBlock}$ which computes a low-rank representation of $\mathbf{L}_{j,s}^D$ given some prescribed rank. The structure of this algorithm is not given here since it is nearly identical to Algorithm 3.

Finally, we turn to the question of how to construct dense diagonal blocks within $\mathbf{L}_j^D$. As in Algorithm 1, lines 7 & 10, we will consider update matrices built from off-diagonal blocks in node j's descendants. In addition, we will need to consider contributions from previously computed low-rank blocks in $\mathbf{L}_j^D$. The details of this process are given in Algorithm 6.

Given algorithms for forming diagonal and off-diagonal blocks in $\mathbf{L}_j^D$, building this matrix follows a straight forward process of iterating over the blocks of $\mathbf{L}_j^D$ in increasing order $s = 1, 2, \ldots$. At each iteration, we either form the diagonal block $\mathbf{L}_{j,s}^D$ or a low-rank decomposition $\mathbf{V}_{j,s}^D (\mathbf{U}_{j,s}^D)^T$. Finally, we note that, unlike the off-diagonal compression scheme

presented in §3.2, forming diagonal blocks and approximate off-diagonals in this order does not guarantee that positive definiteness is maintained throughout the factorization. We discuss our simple method for addressing this issue in §5.3.

4.3. Choosing Diagonal Block Coordinates. In §4.1-4.2 we assumed that indices within $\mathbf{L}_j^D$ could be reordered to expose low-rank structure in off-diagonal blocks $\mathbf{L}_{j,s}^D$. As we discussed in §4.1, this is accomplished by assigning spatial coordinates to degrees of freedom within $\mathbf{L}_j^D$. Indices are reordered such that off-diagonal blocks in $\mathbf{L}_j^D$ describe interactions between spatially separated "pieces" of the separator with which node j is associated. However, we have not yet discussed how these spatial coordinates are determined. In many applications, this information can be determined from the underlying PDE. For example, in §6 we discuss several model problems implemented in the Deal.II finite element library. For these problems, spatial coordinates are determined directly from node positions in a finite element mesh. Unfortunately, this information may not be readily available in some cases. In the interest of building a general, algebraic preconditioner, we wish to also consider cases in which spatial coordinates for system degrees of freedom are not provided.

Research in the area of *graph visualization* has yielded a variety of methods for building visually appealing drawings of graphs [22, 4]. Given a sparse matrix $\mathbf{A}$, we can infer geometric positions for matrix indices by applying these algorithms to the graph structure implied $\mathbf{A}$'s non-zero pattern. In this work, we appeal to *spectral graph drawing* algorithms, which build positions based on the spectral properties of certain matrices associated with the original system matrix $\mathbf{A}$. In particular, we consider $\mathbf{A}$'s *Graph Laplacian* $\mathbb{L}$, defined as follows:

$$\mathbb{L}_{ij} = \begin{cases} -1 & \text{if } i \neq j,\ \mathbf{A}_{ij} \neq 0 \\ 0 & \text{if } i \neq j,\ \mathbf{A}_{ij} = 0 \\ |\{k \neq i : \mathbf{A}_{ik} \neq 0\}| & \text{if } i = j \end{cases} \tag{4.4}$$

We evaluate the three lowest-order eigenvectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ of $\mathbb{L}$ and associate the three-dimensional position $[\mathbf{v}_{1i}\, \mathbf{v}_{2i}\, \mathbf{v}_{3i}]$ with matrix index i. We find that low-accuracy approximations of these eigenvectors suffice, and we evaluate these eigenvectors using Arnoldi iteration. See §6 for further discussion on the cost of constructing these positions.

5. Additional Optimizations and Implementation Details. In this section we discuss additional optimizations for further storage reduction in our algorithm, as well as key implementation details.

5.1. Interior Blocks. The algorithm discussed in §3-4 builds a sparse Cholesky factor on a matrix permuted with a nested dissection ordering. Separators in the nested dissection hierarchy that are sufficiently large – that is, having more than τ_O variables – are identified as supernodes in a supernodal Cholesky factorization and the diagonal and off-diagonal blocks for these supernodes are compressed. Supernodes with fewer than τ_O variables may be factored using the standard supernodal sparse Cholesky algorithm; however, in this section we present a more memory-efficient method for handling these uncompressed blocks.

The collection of compressed separators discussed above partitions the domain in to a collection of mutually disjoint subdomains (see Figure 5.1 – left side), which we refer to as *interior blocks*. This remains true even for non-physical problems in which the "domain" is the graph defined by the sparsity pattern of the matrix to be factored. The nested dissection permutation guarantees that the variables in an interior block appear in a contiguous block in the reordered matrix (see Figure 5.1 – right side). As such, an interior block can be represented by a sequential list of supernode indices. For interior block i (numbered in the order in which it appears in the reordered matrix), we use the notation $\mathbb{B}_i$ to denote the list of supernode indices comprising the block. The column list $\mathcal{C}_i^{\mathbb{B}}$ and off-diagonal row pattern $\mathcal{R}_i^{\mathbb{B}}$

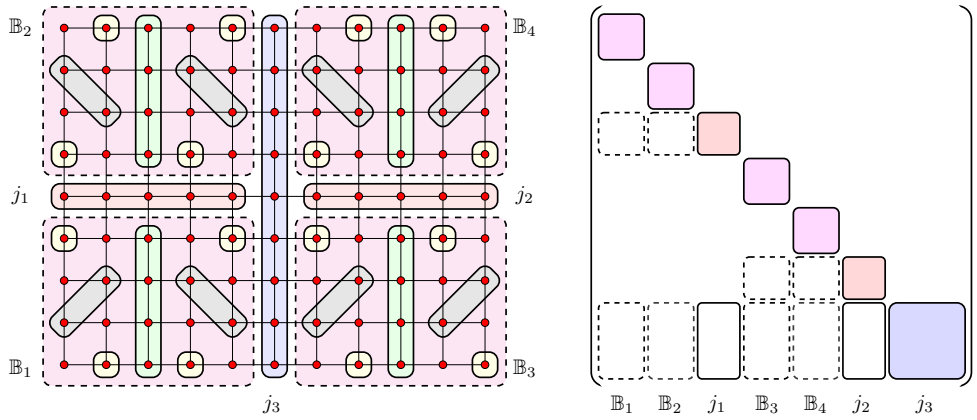

FIG. 5.1. *As in Figure 3.2, we apply nested dissection to a two-dimensional grid and set $\tau_O = 5$. The separators with 5 or more variables partition this domain in to four disjoint subdomains, which we label as interior blocks $\mathbb{B}_{1,\ldots,4}$. These four blocks are shown as regions surrounded by a dashed line in the left image. The block structure in $\mathbf{L}$ resulting from this partitioning is shown on the right. As before, solid white blocks in this matrix denote off-diagonal blocks $\mathbf{L}_j^{\mathcal{O}}$ for which we apply compression according to the methods in §3. Dashed white blocks denote the off-diagonals of interior blocks $\mathbb{B}_{1,\ldots,4}$. These blocks are not stored explicitly. We do, however, store the diagonal components of interior blocks explicitly, and these entries are evaluated using standard, supernodal Cholesky factorization restricted to the interior block subdomain.*

associated with interior block i are defined as follows:

$$\mathcal{C}_i^{\mathbb{B}} = \bigcup_{j \in \mathbb{B}_i} \mathcal{C}_j \qquad \mathcal{R}_i^{\mathbb{B}} = \left(\bigcup_{j \in \mathbb{B}_i} \mathcal{R}_j \right) \setminus \mathcal{C}_i^{\mathbb{B}} \tag{5.1}$$

The matrices $\mathbf{L}(\mathcal{C}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ and $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ are the diagonal and off-diagonal factor blocks for interior block i (the dashed/white and shaded matrix blocks in Figure 5.1, respectively). Indices within interior block i are reordered via nested dissection to guarantee that $\mathbf{L}(\mathcal{C}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ is as sparse as possible. However, for our purposes we can think of interior blocks as representing the "bottom" level of the nested dissection hierarchy. When building a Cholesky factorization, blocks in a nested dissection hierarchy only depend on blocks from lower levels in this hierarchy. Therefore, the factor contents for interior block i are evaluated as follows:

$$\mathbf{L}(\mathcal{C}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}}) = \text{chol}(\mathbf{A}(\mathcal{C}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})) \tag{5.2}$$

$$\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}}) = \mathbf{A}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}}) \mathbf{L}(\mathcal{C}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})^{-T} \tag{5.3}$$

We compute $\mathbf{L}(\mathcal{C}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ using a standard, uncompressed supernodal factorization.

The matrix $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ only stores interactions between the variables of interior block i and supernodes compressed using the methods of §3-4 (see Figure 5.1). As a result, $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ may have many non-zero entries, making it expensive to store explicitly. Fortunately, we can still approximately factor $\mathbf{A}$ without ever explicitly forming the block $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$. In the standard supernodal factorization (Algorithm 1), we explicitly form the Schur complement matrix for each supernode. Here, $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ must be stored explicitly because it is required when running Algorithm 1 on ancestors of nodes in interior block i. The key insight of our approach is that our factorization algorithm only uses $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ in places:

1. We explicitly build Schur complements for *small* diagonal blocks with fewer than τ_D rows/columns. This may depend on contributions from $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$.

2. We must be able to form products of the form $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})\mathbf{G}$ and $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})^T\mathbf{G}$ where $\mathbf{G}$ is an arbitrary dense matrix. These products are required by Algorithms 2 and 5.

Forming the diagonal blocks referred to in the first requirement only necessitates the formation of a sub-block of $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ with at most τ_D rows (see Algorithm 6, line 8). This sub-block is needed exactly once for the formation of a diagonal block. We can form small sub-blocks of $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ as needed to build diagonal block Schur complements, then discard them immediately afterwards.

We also observe that the products from the second requirement listed above can be formed without explicitly forming any part of $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$. Suppose that we wish to compress blocks in supernode j, and that this node has some descendents in interior block i; that is, $\mathbb{D}_j \cap \mathbb{B}_i \neq \emptyset$. Forming the product $\mathbf{L}_j^O\mathbf{G}$ for some dense matrix $\mathbf{G}$ can be done by considering each descendent in $\mathbb{D}_j$ individually, as is done in Algorithm 2. Alternately, we can consider all of the descendents in interior block i simultaneously. To do this, we first recall that $\mathbf{L}_i^{\mathbb{B}} = \mathbf{L}(\mathcal{C}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$ is computed explicitly using sparse supernodal factorization, meaning that its inverse can be applied quickly. It follows from Algorithm 1 that interior block i's contribution to the Schur complement $\boldsymbol{\mathcal{U}}_j^O$ is given by

$$\mathbf{L}(\mathcal{R}_j, \mathcal{C}_i^{\mathbb{B}})\mathbf{L}(\mathcal{C}_j, \mathcal{C}_i^{\mathbb{B}})^T. \tag{5.4}$$

We can use (5.2-5.3) to express the two matrices involved in (5.4) as

$$\mathbf{L}(\mathcal{R}_j, \mathcal{C}_i^{\mathbb{B}}) = \mathbf{A}(\mathcal{R}_j, \mathcal{C}_i^{\mathbb{B}}) \left(\mathbf{L}_i^{\mathbb{B}}\right)^{-T} \qquad \mathbf{L}(\mathcal{C}_j, \mathcal{C}_i^{\mathbb{B}}) = \mathbf{A}(\mathcal{C}_j, \mathcal{C}_i^{\mathbb{B}}) \left(\mathbf{L}_i^{\mathbb{B}}\right)^{-T} \tag{5.5}$$

Finally, we can use (5.5) to write the product of (5.4) and an arbitrary dense matrix $\mathbf{G}$ (as required by Algorithm 2) as

$$\mathbf{L}(\mathcal{R}_j, \mathcal{C}_i^{\mathbb{B}})\mathbf{L}(\mathcal{C}_j, \mathcal{C}_i^{\mathbb{B}})^T\mathbf{G} = \mathbf{A}(\mathcal{R}_j, \mathcal{C}_i^{\mathbb{B}}) \left[\left(\mathbf{L}_i^{\mathbb{B}}\right)^{-T} \left[\left(\mathbf{L}_i^{\mathbb{B}}\right)^{-1} \left[\mathbf{A}(\mathcal{C}_j, \mathcal{C}_i^{\mathbb{B}})^T\mathbf{G}\right]\right]\right] \tag{5.6}$$

As suggested by the parenthesis in (5.6), this product is the result of multiplying a sparse matrix with $\mathbf{G}$, followed by two sparse triangular solves involving $\mathbf{L}_i^{\mathbb{B}}$, followed by another multiplication with a sparse matrix. Since each of these operations can be carried out efficiently, this provides an effective method for forming the matrix products required by Algorithm 2 without having to explicitly store blocks of the form $\mathbf{L}(\mathcal{R}_i^{\mathbb{B}}, \mathcal{C}_i^{\mathbb{B}})$. A similar method can be used to form products with blocks $\mathbf{L}_{j,s}^D$ by replacing $\mathcal{R}_j$ and $\mathcal{C}_j$ in (5.6) (see Algorithm 5) with $\mathcal{R}_{j,s}^D$ and $\mathcal{C}_{j,s}^D$, respectively.

While the optimizations discussed here have the potential to significantly reduce storage requirements, this comes at the cost of somewhat more expensive factorization and triangular solves. We provide concrete examples of this time-memory tradeoff in §6.3.

5.2. Estimating Rank. Up until now, we have assumed when building a low-rank approximations to blocks in $\mathbf{L}$ that the desired rank for each block is known *a priori*. In this section, we discuss how block ranks are chosen. We use our approximate rank-structured Cholesky factor as a preconditioner for the Preconditioned Conjugate Gradient method. Therefore, we are also free to use simple heuristics to determine block ranks, with the understanding that the accuracy with which we approximate blocks will influence the effectiveness of our preconditioner. In principle, we could adaptively approximate blocks up to a certain tolerance (see, e.g., [24, 16]); however, our experiments showed that the additional cost of adaptive approximation outweighed the improved accuracy of preconditioners built with this method. Instead, we use a simple heuristic function depending only on the number of rows and columns in the block to be approximated. In particular, we assign the following rank to a block $\mathbf{B} \in \mathbb{R}^{m\times n}$:

$$\operatorname{rank}(\mathbf{B}) = \alpha\sqrt{k}\log_2(k) + p \tag{5.7}$$

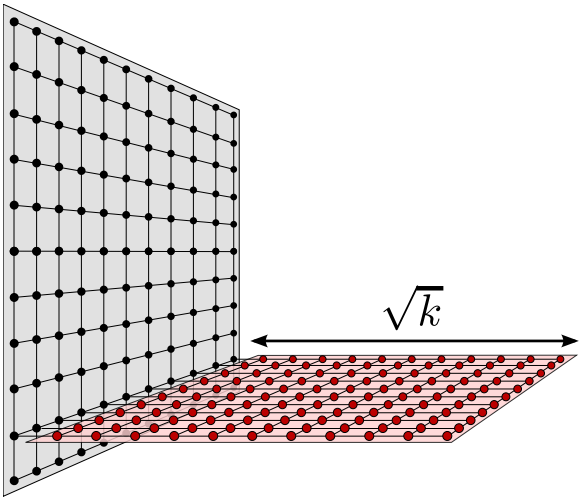

FIG. 5.2. *Interaction between two square separators with grid-based topology. Assuming that each separator has k variables, then the number of variables immediately adjacent to each other in the interaction between these two separators is* $\sqrt{k}$.

where $k = \min(m, n)$ and p is a small oversampling parameter. To provide a brief, intuitive explanation as to why this function was chosen, we consider the interaction between two separators in 3D space. Suppose that the separators take the shape of regular, square, two-dimensional grids intersecting at a right-angle (see Figure 5.2). This may be the case in, for instance, a PDE discretized on a regular, three-dimensional finite difference grid. Assuming that each separator has k variables, there are $\sqrt{k}$ in each separator which are immediately adjacent to the other separator. It is for this reason that we include a term proportional to $\sqrt{k}$ in (5.7). We also scale $\mathrm{rank}(\mathbf{B})$ by $\log_2(k)$ since we empirically observe better preconditioning behavior when larger ranks are used to approximate larger blocks from $\mathbf{L}$. In practice, we also use two different constants in (5.7) – α^D and α^O – which determine ranks during diagonal and off-diagonal compression, respectively.

5.3. Avoiding Indefinite Factorizations. As discussed originally in §4.2, our scheme for compressing diagonal blocks $\mathbf{L}_j^D$ does not provide a guarantee that Schur complements formed during the factorization will remain positive definite. This could result in our factorization algorithm failing for certain inputs. Fortunately, the argument in §3.2 guarantees that in the absence of compression of diagonal blocks $\mathbf{L}_j^D$, all Schur complements remain positive definite. This implies that we can avoid the indefinite Schur complements by approximating diagonal blocks with sufficient accuracy. In practice, we address this issue by adapting the diagonal compression parameter α_D in the event that an indefinite diagonal matrix is encountered during factorization. Specifically, we initialize $\alpha_D = 0.5$ and if factorization fails due to an indefinite matrix, we increase this constant $\alpha_D \longleftarrow 1.25\alpha_D$ and restart the factorization process. This strategy increases the accuracy with which diagonal blocks are approximated until factorization is successful.

6. Results. We have applied the method described in this paper to a number of sample problems. In §6.1 we demonstrate the behavior of our solver on a challenging nonlinear elasticity problem. When applied to the linear systems arising in this problem, our solver provides significant performance improvements over a variety of standard solvers. In §6.2, we discuss the behavior of our solver on a variety of other sample problems. We consider

both standard examples implemented using the *deal.II* finite element analysis library [2, 3] and examples taken from the *University of Florida sparse matrix collection* [8]. While the performance differences between our solver and standard direct and iterative solvers are less dramatic in these examples, these results demonstrate the robustness of our solver.

6.1. An Example: A Nonlinear Elasticity Problem. In this section, we discuss an example that illustrates the behavior of our current solver. While we have tested our solver on numerous problems, the problem described here poses particular difficulty for standard iterative methods. As such, it is an ideal candidate for a hybrid approach such as ours which leverages the reliability of direct solvers with the low memory overhead of iterative methods.

We evaluate our problem using a benchmark problem taken from [27] based on a standard example from the deal.II finite element analysis library [3, 2]. This simulation models quasi-static loading of a nearly-incompressible, hyperelastic block under compression. The code applies a force incrementally over two load steps; at each load step, a nonlinear system of equations is solved to determine the resulting deformation of the block. The nonlinear system is solved by a Newton iteration, and we evaluate the performance of our solver for solving the sequence of linear systems that arise during this process.

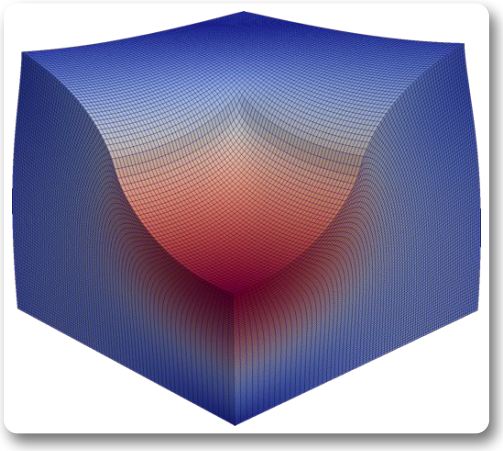

FIG. 6.1. *High-resolution ($N = 80$), elastic block under compression. Our sparse rank-structured preconditioner was used to compute the deformations seen here.*

Because this problem is nearly incompressible, standard displacement-based elements would be prone to locking. Consequently, our test problem uses a mixed formulation with explicit pressure and dilation field variables in addition to the displacement fields. We consider two versions of this problem – henceforth referred to as the $p = 1$ and $p = 2$ problems. In the $p = 1$ problem, displacements are discretized with continuous linear Lagrange brick elements, while pressure and dilation are discretized using discontinuous piecewise constant functions. The $p = 2$ problem discretizes displacements with quadratic elements and uses discontinuous linear elements for pressure and dilation. All variables are discretized on an $N \times N \times N$ element grid. In all problem instances, the pressure and dilation variables are condensed out prior to the linear solve, so the system we solve involves only displacement variables.

6.1.1. Comparison to Standard Iterative Solvers. We have solved the benchmark problem with the preconditioned conjugate gradient (PCG) iteration using our rank-structured Cholesky preconditioner and several other preconditioners provided in Trilinos [17, 20, 21, 19, 18], a library of high-performance solvers developed primarily at Sandia national labs. Our code consistently out-performed a Jacobi preconditioner, an incomplete Cholesky (ICC) preconditioner, and a multi-level (ML) preconditioner in both iteration counts and wall clock time (Figure 6.2(a)–6.2(b)). Our timing results are summarized in Table 6.1.

All results reported in this section were generated on an 8-core Intel Xeon X5570 workstation with 48GB of memory running Ubuntu 12.04, with LAPACK and BLAS implementations provided by the Intel Math Kernel Library version 11.0. We use the preconditioned conjugate gradient (PCG) implementation provided by AztecOO for all tests. All linear systems were solved to a relative ℓ_2 residual error threshold of 10^{-5}. At this accuracy level, the

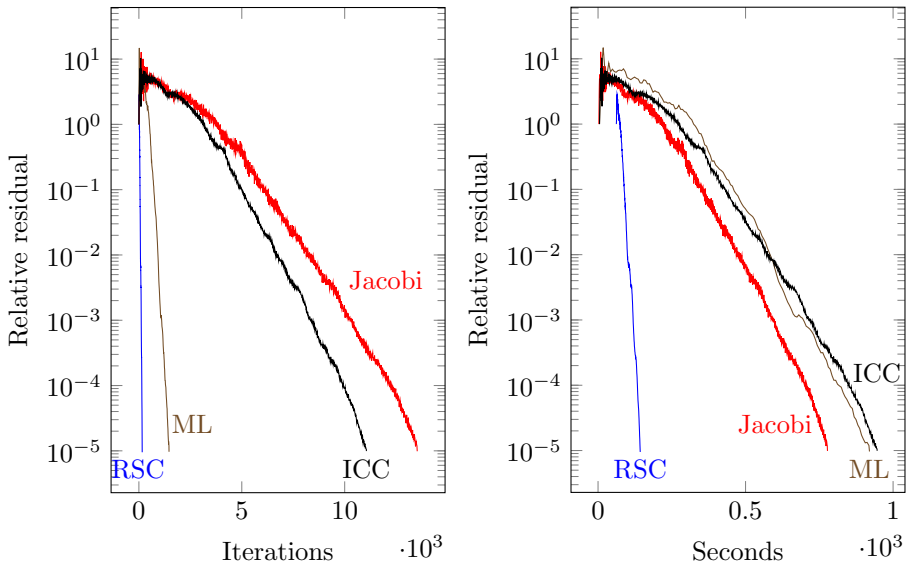

FIG. 6.2. *PCG convergence for the first linear solve in the $p = 1$ benchmark problem with $N = 50$. Jacobi, ICC, ML and RSC refer to solves preconditioned with Jacobi, incomplete Cholesky (IFPACK), multigrid and rank-structured Cholesky (our solver) preconditioners, respectively. Convergence curves relative to wall clock time start at $t > 0$ due to time required to construct the preconditioner.*

nonlinear iteration required 14–15 linear solve steps to converge (as compared to 12 linear solves for a standard Cholesky solver). Due to the time required to solve linear systems using the standard preconditioners, we only ran these example for $N \leq 50$. For the rank-structured Cholesky solver, we ran examples up to $N \leq 80$.

We observe the following properties for the different preconditioners for this problem:

Jacobi: The Jacobi preconditioner (diagonal preconditioner) is simple, but it usually only modestly accelerates convergence. When solving the $p = 1$ problem, more *iterations* are required to converge with the Jacobi preconditioner than with more sophisticated preconditioners. However, each iteration is so cheap that this process requires less *wall clock time* than solves performed with other standard preconditioners. Meanwhile, in the $p = 2$ problem, the number of iterations required when using a Jacobi preconditioner grows considerably, making other standard preconditioners more competitive.

ICC: We timed the PCG iteration using incomplete Cholesky (ICC) preconditioners implemented in both IFPACK and AztecOO. As with most incomplete factorization codes, these solvers require several parameters, including the level of fill allowed in the factorization, drop tolerances dictating which matrix entries should be discarded, and parameters controlling perturbations to the matrix's diagonal. The latter are required to avoid poorly conditioned factorizations since solvers based on incomplete factorizations appear to encounter severe conditioning issues when applied to this problem. In general, "good" parameter choices depend on the problem. We chose parameters for this problem based on experiments with a small problem instance (e.g., $N = 20$ for the $p = 1$ problem). When solving the $p = 1$ problem, even with the tuned parameters, we require almost as many iterations with IFPACK's ICC preconditioner as with the much simpler Jacobi preconditioner. Moreover, the incomplete Cholesky preconditioner costs more than applying the Jacobi preconditioner, so the overall time to solve the linear systems is actually *larger* in this case. Meanwhile, we find that Aztec's ICC preconditioner is unable to make any progress towards convergence in the

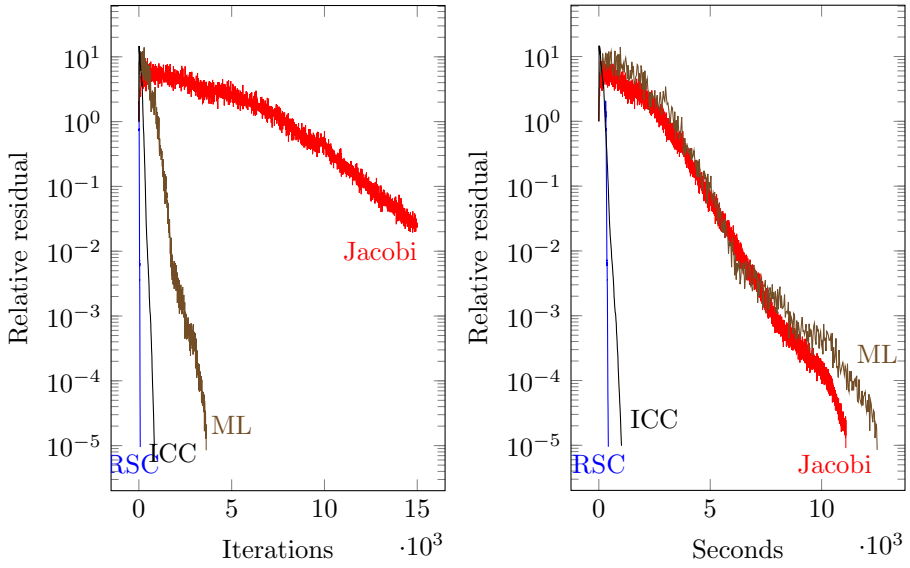

FIG. 6.3. *PCG convergence for the first linear solve in the $p = 2$ benchmark problem with $N = 35$ (see the caption for Figure 6.2(a)-6.2(b)). Note that the Jacobi solve was run to convergence: the left plot is truncated for clarity.*

$p = 1$ problem. We observe the opposite behavior in the $p = 2$ problem. Specifically, IFPACK's ICC preconditioner makes no progress towards convergence, whereas the AztecOO preconditioner performs reasonably well (but still significantly slower than our rank structured solver). In fact, while IFPACK's ICC solver was the least effective standard solver (in terms of wall clock time) applied to the $p = 1$ problem, AztecOO's ICC solver was the *most* effective standard solver applied to the $p = 2$ problem. This phenomenon provides further evidence of the difficulty associated with choosing an effective preconditioner for a given problem.

ML: While multigrid preconditioners perform well on many problems, on our benchmark we see relatively poor convergence and long solve times. We use an algebraic multigrid preconditioner that solves the problem at its coarsest level using a direct solver provided by Amesos. The next coarsest level applies a symmetric Gauss-Seidel smoother over several sweeps (4 in this case). Finer levels use a degree-2 Chebyshev polynomial smoother. These parameters were chosen based on good convergence behavior on a small problem instance (e.g., $N = 20$ for the $p = 1$ problem). When solving the $p = 1$ problem, this solver requires significantly fewer iterations than the Jacobi or incomplete Cholesky solvers; but because applying the preconditioner is relatively expensive, it takes about as long to solve with the multigrid preconditioner as with a Jacobi preconditioner. When applied to the $p = 2$ problem, the multigrid preconditioner requires significantly fewer iterations than the Jacobi solver, but many more than the Aztec00 incomplete Cholesky solver. As a result, this solver is the least effective standard solver that we tested on the $p = 2$ problem. We also note here that multigrid frameworks such as the one provided by ML require tuning a wide variety of parameters (some of which are discussed above). In some cases, tuning problem-specific parameters to achieve good convergence behavior may outweigh the cost of solving the problem with a simpler method.

Rank-Structured Cholesky: The conjugate gradient method preconditioned with our rank-

structured Cholesky solver converges quickly, both in terms of the iteration count and in terms of wall clock time. This is a significant improvement over the other preconditioners.

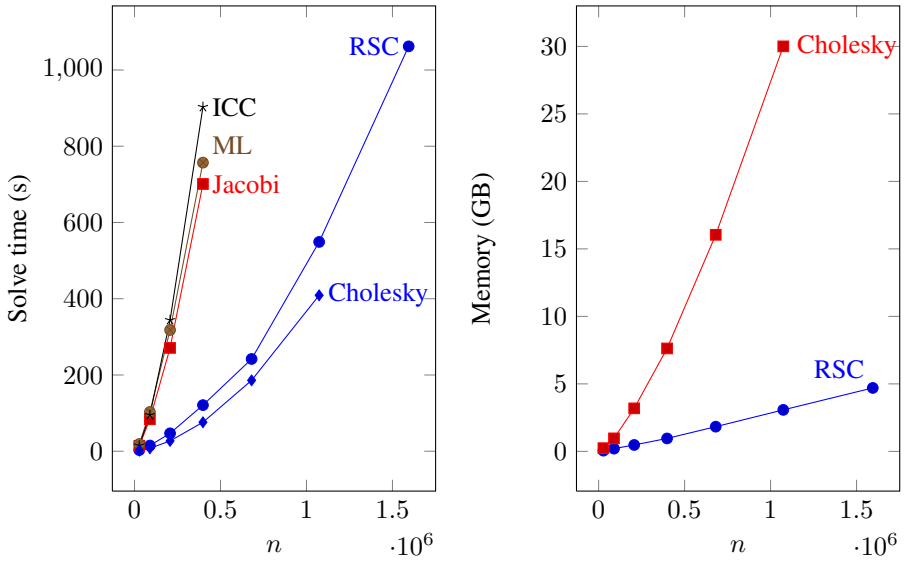

FIG. 6.4. *Mean times and memory usage for the $p = 1$ benchmark problem with problem between $N = 20$ and $N = 80$. Solvers based on direct factorization are much faster than the standard preconditioners in this benchmark. Although it takes somewhat longer to solve systems with our rank-structured solver than with an exact factorization, the memory requirements of the latter approach make it infeasible for larger problems.*

6.1.2. Comparison to Exact Factorization. Beside comparing to standard preconditioners, we also compare our code to an exact sparse Cholesky factorization. As discussed above, our sparse Cholesky implementation closely mirrors CHOLMOD [6] and achieves similar performance and memory usage for exact factorizations. Because CHOLMOD places restrictions on problem size, we use our code for both the rank-structured approximation and the exact Cholesky factorizations

In Figure 6.4, we show how much time and memory we need to solve linear systems with the rank-structured and exact sparse Cholesky factorizations. For this benchmark, both the rank-structured and the exact Cholesky solvers are much faster than the standard preconditioned iterations. The rank-structured Cholesky solver is somewhat slower than the exact Cholesky solver; but the memory requirements of the latter approach make it infeasible for larger problems.

6.2. Other Sample Problems. In this section, we discuss other sample problems arising either from finite element discretizations in Deal.II or from the University of Florida Sparse Matrix collection. For some examples, we also consider the effect of varying the number of power iterations used for low-rank approximation in our solver (see §3.4). We find that increasing this number can result in somewhat more accurate approximation, and a modest reduction in PCG iterations.

Finite element analysis of a trabecular bone: We consider the stiffness matrix produced by finite element analysis of a three-dimensional trabecular bone model. The matrix used

in this problem is provided in the University of Florida sparse matrix collection [1], has dimension $n = 986703$ and has 24419243 non-zeros in it's lower triangular component. We use low-accuracy eigenvectors of the matrix's graph Laplacian to compute three-dimensional coordinates for degrees of freedom in this system (see §4.3). Since no right-hand-side vector is provided for this problem, we solve the system $\mathbf{Ax} = \mathbf{b}$ with the constant vector $\mathbf{b} = (1\ 1\ \ldots\ 1)^T$.

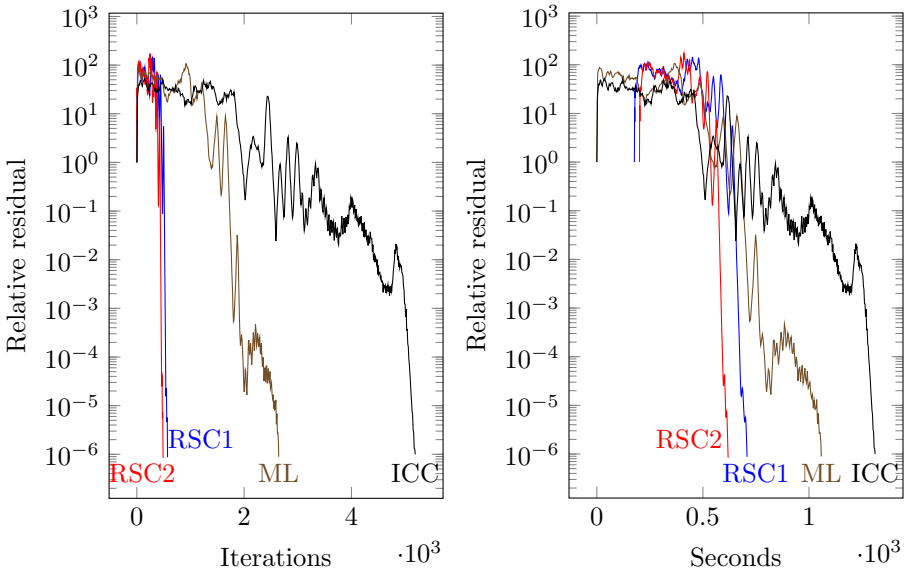

FIG. 6.5. *PCG convergence for the trabecular bone problem. We compare results using incomplete Cholesky (ICC), a multigrid preconditioner (ML), and our solver using $s = 1$ or 2 power iterations when building low-rank approximations (RSC1 and RSC2, respectively). We see that the improved accuracy from more power iterations results in a modest reduction in solution time.*

Finite element analysis of a steel flange: Here, we consider a linear system arising from a three-dimensional mechanical problem discretizing a steel flange. This example can be found in the University of Florida sparse matrix collection [2], has dimension $n = 1564794$ and has 57865083 non-zeros in it's lower triangular component.

Poisson's equation: We apply our solver to Poisson's equation

$$-\nabla \cdot \mathbf{K}(\mathbf{x})\nabla p = f \text{ in } \Omega \tag{6.1}$$

$$p = g \text{ on } \partial\Omega \tag{6.2}$$

where the coefficients $\mathbf{K}(\mathbf{x})$ are optionally both inhomogeneous and anisotropic. The basic setup of this problem follows a standard example from the deal.II library [3]. In particular, rather than solving (6.1-6.2) directly, we define $\mathbf{u} = -\mathbf{K}\nabla p$ and consider the *mixed formu-*

[1] http://www.cise.ufl.edu/research/sparse/matrices/Oberwolfach/bone010.html
[2] http://www.cise.ufl.edu/research/sparse/matrices/Janna/Flan_1565.html
[3] http://www.dealii.org/developer/doxygen/deal.II/step_20.html

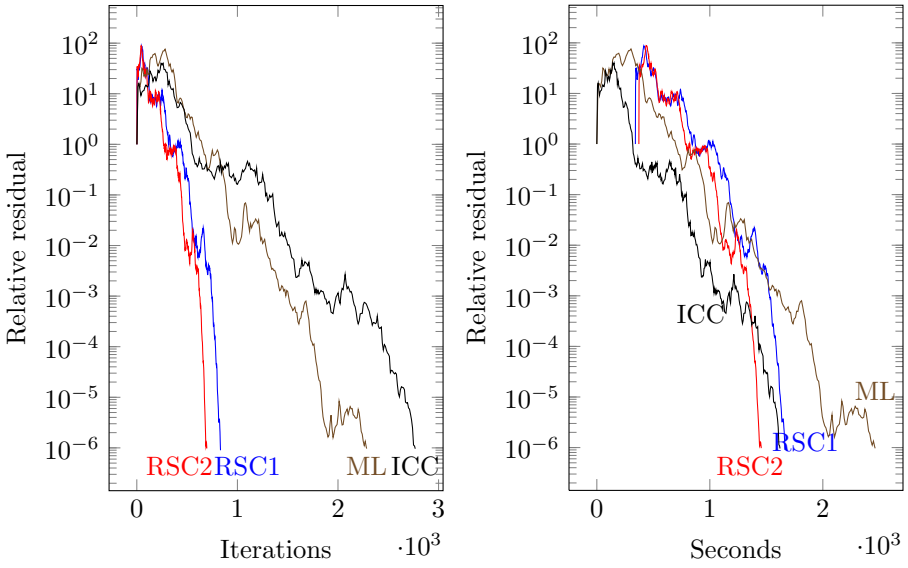

FIG. 6.6. *PCG convergence for the steel flange problem.*

lation of this problem:

$$\mathbf{K}^{-1}\mathbf{u} + \nabla p = 0 \text{ in } \Omega \tag{6.3}$$

$$-\nabla \cdot \mathbf{u} = -f \text{ in } \Omega \tag{6.4}$$

$$p = g \text{ on } \partial\Omega \tag{6.5}$$

This problem is discretized using Raviart-Thomas elements, resulting in a linear system of the form

$$\begin{pmatrix} \mathbf{M} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{U} \\ \mathbf{P} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{g} \end{pmatrix} \tag{6.6}$$

Block elimination of (6.6) yields the following block system:

$$\mathbf{SP} = \mathbf{BM}^{-1}\mathbf{f} - \mathbf{g} \tag{6.7}$$

$$\mathbf{MU} = \mathbf{f} - \mathbf{B}^T\mathbf{P} \tag{6.8}$$

where $\mathbf{S}$ is the positive definite Schur complement matrix $\mathbf{S} = \mathbf{BM}^{-1}\mathbf{B}$. (6.7) can be solved via PCG given an efficient procedure for forming matrix-vector products with $\mathbf{S}$. This in turn requires efficient application of $\mathbf{M}^{-1}$. This will also be accomplished with the conjugate gradient method. In the case $\mathbf{K}(\mathbf{x}) = const$, linear systems involving $\mathbf{M}$ turn out to be quite easy to solve. In fact, these systems can be solved quickly using standard conjugate gradients with no preconditioning. Nevertheless, this problem provides a useful benchmark for our solver. We also consider versions of the problem in which $\mathbf{K}(\mathbf{x})$ is anisotropic and highly inhomogeneous to demonstrate that our solver can also handle these cases.

6.3. Comparisons. In this section we provide timing and memory usage statistics for some of the features discussed in §3-5.

Diagonal block coordinates: In §4.3 we discussed how spatial coordinates for diagonal block indices can be chosen either geometrically based on information obtained directly from a problem's discretization, or algebraically via approximations of the low-order eigenvectors associated with the problem's graph Laplacian. We compare both approaches applied to a linear system taken from the nonlinear elasticity benchmark problem (§6.1). We find that both approaches produce preconditioners which converge in a comparable number of iterations. Recall that these spatial coordinates are used to permute indices within compressed diagonal blocks. We also consider results when randomly permuting these indices to demonstrate the need for an effective permutation. In this case we observe a significant increase in the number of required PCG iterations. We also note that the diagonal rank constant τ_D had to be increased several times in this case to allow for successful factorization (see §5.3). Convergence plots for these three diagonal block orderings are provided in Figure 6.7. Finally, it is worth noting that the ordering provided by the original problem discretization and nested dissection ordering may be sufficient for diagonal block compression. For the problem considered here we find that we observe little difference in the PCG convergence behavior even when no additional permutations are applied to diagonal block indices. However, it should be noted that this is not guaranteed to be the case, and a poor choice of diagonal coordinates can lead to poor performance (see the random reordering result in Figure 6.7).

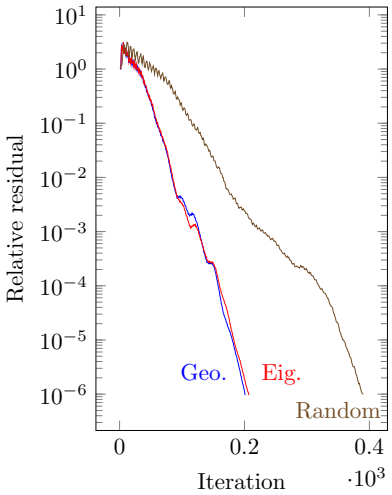

FIG. 6.7. *PCG convergence using random diagonal block reordering, as well as reordering based on geometric index positions (Geo.) and positions obtained from a low-accuracy eigensolve (Eig.).*

Interior block performance: In §5.1 we discussed how to avoid building and storing certain factor blocks explicitly to further reduce memory usage. While this method reduces storage requirements, it comes at the cost of increased factorization and triangular solve time. Here we compare memory usage and solution time results for two examples computed with and without this optimization.

7. Conclusions. In this paper, we have described a direct factorization method for the solution of large sparse linear systems that arise from PDE discretizations. Like standard direct solvers, our approach is black box, and can work with the pre-assembled matrix without prior information about details of an underlying mesh or a specific discretization method. By taking advantage of the low-rank block structure that arises from the underlying PDE, our method requires significantly less memory than standard direct methods, but through careful code organization we retain the high performance of standard direct solvers through use of level-3 BLAS and LAPACK calls. We have demonstrated through examples that our approach retains much of the robustness of standard direct solvers, and yields a faster time to solution than the standard multilevel algebraic multigrid preconditioner ML.

Limitations and Future Work: So far, our work has focused solely on symmetric and positive definite matrices. Other authors have showed how to deal with indefinite problems, with a particular focus on Helmholtz equations, and we intend to adapt that work along with standard static pivoting approaches developed in the context of ordinary sparse parallel LU decompositions. We also so far only have limited parallelism through threaded BLAS calls, but intend to extend our code to work in a distributed memory setting in the future.

REFERENCES

[1] SATISH BALAY, JED BROWN, KRIS BUSCHELMAN, VICTOR EIJKHOUT, WILLIAM D. GROPP, DINESH KAUSHIK, MATTHEW G. KNEPLEY, LOIS CURFMAN MCINNES, BARRY F. SMITH, AND HONG ZHANG, *PETSc users manual*, Tech. Report ANL-95/11 - Revision 3.3, Argonne National Laboratory, 2012.

[2] W. BANGERTH, R. HARTMANN, AND G. KANSCHAT, *deal.II – a general purpose object oriented finite element library*, ACM Trans. Math. Softw., 33 (2007), pp. 24/1–24/27.

[3] W. BANGERTH, T. HEISTER, AND G. KANSCHAT, `deal.II` *Differential Equations Analysis Library, Technical Reference*. `http://www.dealii.org`.

[4] GIUSEPPE DI BATTISTA, PETER EADES, ROBERTO TAMASSIA, AND IOANNIS G. TOLLIS, *Graph Drawing: Algorithms for the Visualization of Graphs*, Prentice Hall, 1998.

[5] S. CHANDRASEKARAN, P. DEWILDE, M. GU, AND N. SOMASUNDERAM, *On the numerical rank of the off-diagonal blocks of schur complements of discretized elliptic PDEs*, SIAM Journal on Matrix Analysis and Applications, 31 (2010), pp. 2261–2290.

[6] YANQING CHEN, TIMOTHY A. DAVIS, WILLIAM W. HAGAR, AND SIVASANKARAN RAJAMANICKAM, *Algorithm 887: CHOLMOD, supernodal sparse cholesky factorization and update/downdate*, ACM Transactions on Mathematical Software, 35 (2008).

[7] TIMOTHY A. DAVIS, *Direct Methods for Sparse Linear Systems*, Society for Industrial and Applied Mathematics, 2006.

[8] TIMOTHY A. DAVIS AND YIFAN HU, *The university of florida sparse matrix collection*, ACM Trans. Math. Softw., 38 (2011), pp. 1:1–1:25.

[9] I. S. DUFF, A. M. ERISMAN, AND J. K. REID, *Direct Methods for Sparse Matrices*, Oxford University Press, 1986.

[10] BJÖRN ENGQUIST AND LEXING YING, *Sweeping preconditioner for the helmholtz equation: Hierarchical matrix representation*, Communications on Pure and Applied Mathematics, 64 (2011), pp. 697–735.

[11] ALAN GEORGE, *Nested dissection of a regular finite element mesh*, SIAM Journal on Numerical Analysis, 10 (1973).

[12] A. GEORGE AND J. W. H. LIU, *Computer Solution of Large Sparse Positive Definite Systems*, Prentice-Hall, 1986.

[13] A GILLMAN AND PG MARTINSSON, *A direct solver with o(n) complexity for variable coefficient elliptic pdes discretized via a high-order composite spectral collocation method*, SIAM Journal on Scientific Computing, 36 (2014), pp. A2023–A2046.

[14] LARS GRASEDYCK, RONALD KRIEMANN, AND SABINE LE BORNE, *Parallel black box\ mathcal {H}-lu preconditioning for elliptic boundary value problems*, Computing and Visualization in Science, 11 (2008), pp. 273–291.

[15] LESLIE GREENGARD, DENIS GUEYFFIER, AND PER-GUNNAR MARTINSSON VLADIMIR ROKHLIN, *Fast direct solvers for integral equations in complex three-dimensional domains*, Acta Numerica, 18 (2009), pp. 243–275.

[16] N. HALKO, P. G. MARTINSSON, AND J. A. TROPP, *Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions*, SIAM Review, 53 (2011), pp. 217–288.

[17] MICHAEL HEROUX, ROSCOE BARTLETT, VICKI HOWLE ROBERT HOEKSTRA, JONATHAN HU, TAMARA KOLDA, RICHARD LEHOUCQ, KEVIN LONG, ROGER PAWLOWSKI, ERIC PHIPPS, ANDREW SALINGER, HEIDI THORNQUIST, RAY TUMINARO, JAMES WILLENBRING, AND ALAN WILLIAMS, *An Overview of Trilinos*, Tech. Report SAND2003-2927, Sandia National Laboratories, 2003.

[18] MICHAEL A. HEROUX, ROSCOE A. BARTLETT, VICKI E. HOWLE, ROBERT J. HOEKSTRA, JONATHAN J. HU, TAMARA G. KOLDA, RICHARD B. LEHOUCQ, KEVIN R. LONG, ROGER P. PAWLOWSKI, ERIC T. PHIPPS, ANDREW G. SALINGER, HEIDI K. THORNQUIST, RAY S. TUMINARO, JAMES M. WILLENBRING, ALAN WILLIAMS, AND KENDALL S. STANLEY, *An overview of the trilinos project*, ACM Trans. Math. Softw., 31 (2005), pp. 397–423.

[19] MICHAEL A. HEROUX AND JAMES M. WILLENBRING, *Trilinos Users Guide*, Tech. Report SAND2003-2952, Sandia National Laboratories, 2003.

[20] MICHAEL A. HEROUX, JAMES M. WILLENBRING, AND ROBERT HEAPHY, *Trilinos Developers Guide*, Tech. Report SAND2003-1898, Sandia National Laboratories, 2003.

[21] ———, *Trilinos Developers Guide Part II: ASCI Software Quality Engineering Practices Version 1.0*, Tech. Report SAND2003-1899, Sandia National Laboratories, 2003.

[22] MICHAEL KAUFMANN AND DOROTHEA WAGNER, eds., *Drawing graphs: methods and models*, Springer-Verlag, London, UK, UK, 2001.

[23] S. LI, M. GU, C. WU, AND JIANLIN XIA, *New efficient and robust hss cholesky factorization of spd matrices*, SIAM Journal on Matrix Analysis and Applications (to appear), 33 (2012), pp. 886–904.

[24] EDO LIBERTY, FRANCO WOOLFE, PER-GUNNAR MARTINSSON, VLADIMIR ROKHLIN, AND MARK TYGERT, *Finding structure with randomness: Probabilistic algorithms for constructing approximate*

matrix decompositions, Proceedings of the National Academy of Science, 104 (2007), pp. 20167–20172.

[25] PER-GUNNAR MARTINSSON, *A fast direct solver for a class of elliptic partial differntial equations*, Journal of Scientific Computing, 38 (2009), pp. 316–330.

[26] GENE POOLE, YONG-CHENG LIU, AND JAN MANDEL, *Advancing analysis capabilities in ansys through solver technology*, Electronic Transactions on Numerical Analysis, 15 (2003), pp. 106–121.

[27] S. REESE, P. WRIGGERS, AND B. D. REDDY, *A new locking-free brick element technique for large deformation problems in elasticity*, Computers and Structures, 75 (2000), pp. 291–304.

[28] PHILLIP G SCHMITZ AND LEXING YING, *A fast nested dissection solver for cartesian 3d elliptic problems using hierarchical matrices*, Journal of Computational Physics, 258 (2014), pp. 227–245.

[29] JIANLIN XIA, *Randomized sparse direct solvers*, SIAM Journal on Matrix Analysis and Applications (to appear), (2012).

[30] JIANLIN XIA, SHIVKUMAR CHANDRASEKARAN, MING GU, AND XIAOYE S. LI, *Superfast multifrontal method for large structured systems of equations*, SIAM Journal on Matrix Analysis and Applications, 31 (2009).

Algorithm 5: diagonalMultiply: Given an input matrix $\mathbf{G}$, forms the product $\mathbf{L}_{j,s}^D\mathbf{G}$, where $\mathbf{L}_{j,s}^D$ is an off-diagonal block in $\mathbf{L}_j^D$. This is done without forming $\mathbf{L}_{j,s}^D$ explicitly. The alignSet function performs the following operation: $\text{alignSet}(S_1, S_2) = \{i - \min(S_2) + 1 : i \in S_1\}$. We will use this function to express row and column sets for block s relative to other blocks in the hierarchy.

input : j, $\mathbf{G}$, transpose, s
output: $\mathbf{V}_{j,s}^D, \mathbf{U}_{j,s}^D$

begin
  Apply necessary diagonal block inverse to $\mathbf{G}$
  $s_1 =$ left child of s
  $\mathbf{G} \longleftarrow \text{diagonalSolve}(j, \mathbf{G}, \text{transpose} = \text{true}, s_1)$
  // Initialize a workspace for multiplication
  $\mathbf{W} \longleftarrow \mathbf{A}(\mathcal{R}_{j,s}^D, \mathcal{C}_{j,s}^D)\mathbf{G}$
  // Apply contributions from supernode descendants
  for each $k \in \mathbb{D}_j$ **do**
    // Extract the required sub-matrix from $\mathbf{G}$
    $\mathbf{G}_{sub} \longleftarrow \text{gatherRows}(\mathbf{G}, \mathcal{C}_{j,s}^D, \mathcal{R}_{k\to j}^D \cap \mathcal{C}_{j,s}^D)$
    // Form the needed product with two multiplications
    $\mathbf{T} \longleftarrow \left[\mathbf{L}_k^O(R_{k\to j}^D \cap C_{j,s}^D, :)\right]^T \mathbf{G}_{sub}$
    $\mathbf{T} \longleftarrow \mathbf{L}_k^O(R_{k\to j}^D \cap R_{j,s}^D, :)\mathbf{T}$
    // Scatter result to the output matrix
    $\mathbf{W} \longleftarrow \mathbf{W} - \text{scatterRows}(\mathbf{T}, \mathcal{R}_{k\to j}^D \cap \mathcal{R}_{j,s}^D, \mathcal{R}_{j,s}^D)$

  // Apply contributions from other blocks in $\mathbf{L}_j^D$
  $p =$ parent of s
  while p is not null **do**
    // Block p contributes to s only if it appears
    // earlier in the ordering
    if $p < s$ **then**
      // Get necessary row and column ranges from p
      // This is valid since $R_{j,s}^D \subset R_{j,p}^D$ and $C_{j,s}^D \subset R_{j,p}^D$
      $R_{sub} = \text{alignSet}(R_{j,s}^D, R_{j,p}^D)$
      $C_{sub} = \text{alignSet}(C_{j,s}^D, R_{j,p}^D)$
      // Perform multiplication similar to (3.9)
      $\mathbf{U}_{prod} \longleftarrow (\mathbf{U}_{j,p}^D)^T\mathbf{U}_{j,p}^D$
      $\mathbf{T} \longleftarrow [\mathbf{V}_{j,p}^D(C_{sub}, :)]^T\mathbf{G}$
      $\mathbf{T} \longleftarrow \mathbf{U}_{prod}\mathbf{T}$
      $\mathbf{T} \longleftarrow \mathbf{V}_{j,p}^D(R_{sub}, :)\mathbf{T}$
      Accumulate result in workspace
      $\mathbf{W} \longleftarrow \mathbf{W} - \mathbf{T}$
    Continue moving up the block hierarchy
    $p \longleftarrow$ parent of p
  return $\mathbf{W}$

Algorithm 6: factorDiagonal: Builds a factored diagonal block $\mathbf{L}_{j,s}^D$ within $\mathbf{L}_j^D$. s is assumed to be the index of a diagonal block. We introduce two new pieces of notation here: $R_{k\to j,s}^D = \{1 \le p \le |\mathcal{R}_k| : r_k^p \in \mathcal{C}_{j,s}^D\}$ and $\mathcal{R}_{k\to j,s}^D = \{r_k^p \in \mathcal{R}_k : r_k^p \in \mathcal{C}_{j,s}^D\}$. That is, these are rows from descendant k which are relevant to the formation of diagonal block s within $\mathbf{L}_j^D$.

input : j, s
output: $\mathbf{L}_{j,s}^D$

begin
  // Initialize Schur complement with matrix contents.
  // This is a diagonal block, so $\mathcal{R}_{j,s}^D = \mathcal{C}_{j,s}^D$.
  $\mathcal{U}_{j,s}^D \longleftarrow \mathbf{A}(\mathcal{R}_{j,s}^D, \mathcal{C}_{j,s}^D)$
  // Accumulate contributions from descendants.
  for each $k \in \mathbb{D}_j$ **do**
    // Build dense update block
    diagUpdate $\longleftarrow \mathbf{L}_k^D(R_{k\to j,s}^D, :)\mathbf{L}_k^D(R_{k\to j,s}^D, :)^T$
    // Scatter updates to the Schur complement
    $\mathcal{U}_{j,s}^D \longleftarrow \mathcal{U}_{j,s}^D - \text{scatter}(\text{diagUpdate}, \mathcal{R}_{k\to j,s}^D, \mathcal{C}_{j,s}^D, \mathcal{R}_{k\to j,s}^D, \mathcal{C}_{j,s}^D)$

  // Accumulate contributions from previous blocks
  // in $\mathbf{L}_j^D$.
  $p \longleftarrow$ parent of s
  while p is not null **do**
    // Block p contributes to s only if it appears earlier
    // in the ordering
    if $p < s$ **then**
      // Get necessary row and column ranges from p
      // This is valid since $R_{j,s}^D \subset R_{j,p}^D$ and $C_{j,s}^D \subset R_{j,p}^D$
      $R_{sub} = \text{alignSet}(R_{j,s}^D, R_{j,p}^D)$
      $C_{sub} = \text{alignSet}(C_{j,s}^D, R_{j,p}^D)$
      // Build a dense update matrix
      $\mathbf{U}_{prod} \longleftarrow (\mathbf{U}_{j,p}^D)^T\mathbf{U}_{j,p}^D$
      $\mathbf{T} \longleftarrow \mathbf{U}_{prod}[\mathbf{V}_{j,p}^D(C_{sub}, :)]^T$
      $\mathbf{T} \longleftarrow \mathbf{V}_{j,p}^D(R_{sub}, :)\mathbf{T}$
      // Subtract update from Schur complement
      $\mathcal{U}_{j,s}^D \longleftarrow \mathcal{U}_{j,s}^D - \mathbf{T}$
  // Factor node j's diagonal block
  $\mathbf{L}_{j,s}^D \longleftarrow \texttt{Cholesky}(\mathcal{U}_{j,s}^D)$
  return $\mathbf{L}_{j,s}^D$

	N	20	30	40	50	60	70	80
	n	27783	89373	206763	397953	680943	1073733	1594323
Jacobi	Total time (s)	194	1170	4065	10520			
	Mean time (s)	14	84	271	701			
	Total iterations	49579	87957	133361	170750			
	Mean iterations	3541	6282	8890	11383			
ICC (IFPACK)	Total time (s)	207	1334	5159	13540			
	Mean time (s)	15	95	344	903			
	Total iterations	48637	87244	132981	168172			
	Mean iterations	3474	6231	8865	12110			
ML	Total time (s)	270	1447	4763	11360			
	Mean time (s)	19	103	318	757			
	Total iterations	5949	10170	15249	19109			
	Mean iterations	425	726	1017	1274			
RSC	Total time (s)	32	173	618	1470	2862	6174	12350
	Mean time (s)	2.3	12	41	98	191	412	823
	Total iterations	349	615	1209	1514	1729	2278	4050
	Mean iterations	24	41	80	100	115	151	270

TABLE 6.1

Nonlinear elasticity ($p = 1$) performance results: *Relative performance of standard iterative methods (Jacobi, incomplete Cholesky, and multigrid) compared to rank-structured Cholesky.*

	N	10	15	20	25	30	35	40	45
	n	27783	89373	206763	397953	680943	1073733	1594323	
Jacobi	Total time (s)	501	3731	13610	33920				
	Mean time (s)	36	267	907	2261				
	Total iterations	60289	110971	175525	231833				
	Mean iterations	4306	7926	11701	15455				
ICC (Aztec)	Total time (s)	61	287	955	2473	5097	10150		
	Mean time (s)	4.4	20	64	165	340	677		
	Total iterations	897	1590	2689	3914	5341	7030		
	Mean iterations	64	113	179	260	356	468		
ML	Total time (s)	918	4431	13920	34950	73690	144900	227800	
	Mean time (s)	66	317	928	2330	4913	9660	15187	
	Total iterations	9149	13770	20748	26534	32292	40126	43509	
	Mean iterations	653	983	1383	1768	2152	2675	2900	
RSC	Total time (s)	42	200	616	1460	2844	5139	9873	14900
	Mean time (s)	3	13	41	97	190	668	658	993
	Total iterations	121	196	343	438	497	668	1438	1429
	Mean iterations	8	13	22	29	33	44	95	95

TABLE 6.2

Nonlinear elasticity ($p = 2$) performance results: *Relative performance of standard iterative methods (Jacobi, incomplete Cholesky, and multigrid) compared to rank-structured Cholesky.*

Symbol	Meaning
$\mathbf{A}$	System matrix (generally assumed to have been permuted with fill-reducing ordering).
$\mathbf{L}$	Cholesky factor matrix (of permuted system).
$c_{j,s}, c_{j,e}$	First and last column in supernode j.
$\mathcal{C}_j$	Supernode j's column set.
$\mathcal{C}_j^O$	Set of columns occurring after supernode j.
$\mathcal{R}_j = \{r_j^1, r_j^2, \ldots\}$	Set of non-zero rows in supernode j's off-diagonal.
$\mathbf{L}_j$	Supernode j's block column in $\mathbf{L}$.
$\mathbf{L}_j^D, \mathbf{L}_j^O$	Diagonal and off-diagonal blocks of $\mathbf{L}_j$, respectively. $\mathbf{L}_j^D \in \mathbb{R}^{\lvert\mathcal{C}_j\rvert \times \lvert\mathcal{C}_j\rvert}$ and $\mathbf{L}_j^O \in \mathbb{R}^{\lvert\mathcal{R}_j\rvert \times \lvert\mathcal{C}_j\rvert}$.
$\mathfrak{U}_j^D, \mathfrak{U}_j^D$	Schur complement matrices corresponding to $\mathbf{L}_j^D$ and $\mathbf{L}_j^O$
$\mathbb{D}_j$	Indices of supernode descendants of supernode j; $\mathbb{D}_j = \{1 \leq k < j : \mathcal{R}_k \cap \mathcal{C}_j \neq \emptyset\}$.
$R_{k\to j}^D, R_{k\to j}^O$	Recalling that $\mathcal{R}_k = \{r_k^1, r_k^2, \ldots\}$, $R_{k\to j}^D = \{1 \leq p \leq \lvert\mathcal{R}_k\rvert : r_k^p \in \mathcal{C}_j\}$ and $R_{k\to j}^O = \{1 \leq p \leq \lvert\mathcal{R}_k\rvert : r_k^p \in \mathcal{R}_j\}$. That is the set of rows in node k needed to construct $\mathbf{L}_j^D$ and $\mathbf{L}_j^O$, respectively (relative to the row space of $\mathbf{L}_k^O$).
$\mathcal{R}_{k\to j}^D, \mathcal{R}_{k\to j}^O$	Similar to the definitions above, $\mathcal{R}_{k\to j}^D = \{r_k^p \in \mathcal{R}_k : r_k^p \in \mathcal{C}_j\}$ and $\mathcal{R}_{k\to j}^O = \{r_k^p \in \mathcal{R}_k : r_k^p \in \mathcal{R}_j\}$. That is the set of rows in node k needed to construct $\mathbf{L}_j^D$ and $\mathbf{L}_j^O$, respectively (relative to the full row space of $\mathbf{L}$).
$R_{k\to j,s}^D, \mathcal{R}_{k\to j,s}^D$	Similar to $R_{k\to j}^D$ and $\mathcal{R}_{k\to j}^D$. $R_{k\to j,s}^D = \{1 \leq p \leq \lvert\mathcal{R}_k\rvert : r_k^p \in \mathcal{C}_{j,s}^D\}$ and $\mathcal{R}_{k\to j,s}^D = \{r_k^p \in \mathcal{R}_k : r_k^p \in \mathcal{C}_{j,s}^D\}$. That is, these are rows from descendant k which are relevant to the formation of diagonal block s within $\mathbf{L}_j^D$.
$\mathbf{V}_j, \mathbf{U}_j$	Low-rank representation for supernode j's off-diagonal block ($\mathbf{L}_j^O \approx \mathbf{V}_j \mathbf{U}_j^T$).
$\mathbf{L}_{j,s}^D$	Block s from the diagonal matrix $\mathbf{L}_j^D$.
$\mathbf{V}_{j,s}^D, \mathbf{U}_{j,s}^D$	Low-rank representation for block s of $\mathbf{L}_j^D$.
$R_{j,s}^D, C_{j,s}^D$	Row and column sets over which $\mathbf{L}_{j,s}^D$ is defined, relative to the dense matrix $\mathbf{L}_j^D$.
$\mathcal{R}_{j,s}^D, \mathcal{C}_{j,s}^D$	Row and column sets over which $\mathbf{L}_{j,s}^D$ is defined, relative to the full system $\mathbf{A}$ (or $\mathbf{L}$).

TABLE 7.1
List of symbols